A Channel Theoretic Approach to Conditional Reasoning

Lawrence Cavedon

Doctor of Philosophy
University of Edinburgh
1995

Declaration

I declare that this thesis has been composed by myself and that the research reported here has been conducted by myself unless otherwise indicated.

Lawrence Cavedon

Edinburgh, May 1 1995

Abstract

Channel Theory is a recently developed mathematical model of information flow, based on ideas emanating from situation theory. Channel theory addresses a number of important properties of information flow, such as context-dependence, modularity of information, and the possibility of error. This thesis is concerned with the use of channel theory as a formal framework for various constructs relating to conditional sentences. In particular, the main concern is to obtain logics for reasoning about conditionals, generics and default properties within the channel theoretic framework.

After presenting the philosophical and mathematical foundations of channel theory, the viability of using this framework as a basis for a logic of conditionals is examined. It is demonstrated that the channel operations defined by Barwise and Seligman are unsuitable for this task since they support certain unacceptable instances of *transitivity* and *monotonicity*. The problem seems to be that, while the contextual nature of a channel *reflects* the implicit background conditions of a regularity, these conditions are not adequately represented in such a way as to account for them in the operations themselves.

A method is proposed for implicitly representing the background assumptions of a channel, by way of a *subchannel relation,* which results in an ordering being defined over any given collection of channels. As required by the maxims of situated reasoning, the background assumptions of a particular channel are not represented within that channel itself, but are instead distributed throughout an ordered collection of channels. The channel operations are modified so as to account for background assumptions encoded in this way, leading to a channel theoretic logic of conditionals that invalidates the unacceptable rules of inference. It is shown that by imposing certain simple, independently motivated conditions on channel hierarchies, a powerful logic of conditionals, supporting many desirable patterns of reasoning, is obtained.

A channel theoretic analysis of generic sentences is proposed and is shown to have several useful properties. In particular, the contextual component provided by a channel addresses certain examples that prove problematic to normative accounts of generics. By making a simple assumption—that any given collection of channels is self-contained with respect to implicit background assumptions—a channel theoretic logic of generics is obtained directly from the previous logic of conditionals. This logic is shown to support several important patterns of reasoning, including the principle of *specificity*.

In order to obtain a system for defeasibly reasoning about individuals from the channel theoretic logic of generics, a *maximal normality* condition is defined. This condition reflects an agent's assumption that individuals are "normal" (with respect to some collection of properties) unless that assumption conflicts with other information held by the agent. The resulting system is shown to effectively define a *preferential consequence relation*, in the sense of Kraus, Lehmann and Magidor. The channel theoretic framework suggests an approach to defeasible reasoning that involves a shift in methodology from standard approaches and seems aligned to the particular concerns of situated reasoning. Basically, rather than viewing defeasible reasoning as reasoning with incomplete information, the channel theoretic approach is to view the problem as reasoning with approximate regularities which hold only within some limited context. This methodology is illustrated via an application to the *qualification problem* from the AI planning literature.

The thesis concludes with a brief discussion of the potential use of channel theory as a mathematical framework for situated reasoning. In particular, the indexical nature of reasoning in a multi-agent environment seems particularly relevant to the particular issues addressed by channel theory.

Acknowledgements

Firstly and foremostly I would like to thank my supervisor Robin Cooper for listening, providing much useful feedback, and especially for showing patience and encouragement during the seemingly many moments of self-doubt.

I feel enormously privileged to have been able to study at the Centre for Cognitive Science and I find it hard to imagine a better environment in which to be a research student—I'd like to thank all students and staff (academic, technical and administrative) for their contributions to the special feel of the place. A number of people need to be singled out for special thanks, however. Paul Schweizer (my second supervisor) provided lots of useful guidance and encouragement in my early days in Edinburgh. Pat Healey, Martin Mellor, Rodger Kibble and Lex Holt were office mates at various stages of the business end of my PhD—I had a great time sharing space with them and also learnt heaps on the way (I'm sure I also owe each of them apologies, especially to Pat for my thick-headedness and to Rodger and Marts for continually drumming on the desk). Former Edinburgher Jerry Seligman (who along with Jon Barwise developed the framework on which this thesis is based) provided lots of useful feedback on the work described in Chapter 3. Sheila Glasbey collaborated on aspects of the work described in Chapter 4 and also read a draft of that chapter. I had lots of interesting discussions with various others at various stages on various topics related to my PhD, especially Carl Vogel, Peter Ruhrberg (even though he never seemed to agree with anything I said!), Andreas Schöter, Matthias Paul, Jonathon Ginzburg, Joris Hulstijn and other members of the Meaning and Computation group. Finally, Rosanna Maccagnano helped re-establish the link with Robin after I moved back to Melbourne and Betty Hughes provided for all my administrative needs both there and here.

Apart from the particular academic benefits, it was a joy to be able to spend three years in Edinburgh. Lots of people helped make that possible. Thanks especially to Ed Uni Korfball Club for lots of good friends and fun times (and a British championship!). Thanks also to Sue and Bill for providing a roof over my head (and lots of yummy meals and good company) during the final few months there.

Financially, I owe a debt to The British Council who (via the administration of the Association of Commonwealth Universities) funded my time in Edinburgh and thereby made it possible. Since returning to Melbourne, I have been employed at the Royal Melbourne Institute of Technology, which has put food on the table and also given me a forum in which to continue research within the AI group there. The University of Melbourne accommodated me as a visiting researcher between the time of my arrival in Melbourne and the commencement of my current employment.

Lastly and mostly, I have to thank my wife Barb who has been through it all with me, (temporarily!) interrupting her own career, listening to me, proofreading, believing in me even when I didn't, and generally sharing the ride. Without her I probably never would have started, I don't know how I could have got through the end, and the middle bits would have been no fun at all.

Chapter 1

Introduction

This thesis is concerned with the application of Barwise and Seligman's *Channel Theory* (Barwise and Seligman 1993; Barwise and Seligman 1994; Seligman 1993; Seligman and Barwise 1993), a framework for modelling information flow that has its roots in Barwise and Perry's (1983) *Situation Theory*. In particular, I investigate the use of channel theory as a semantic framework for reasoning with conditionals, generics and default rules. These various tasks have much in common and using a unified framework to model each of them offers the possibility of insights into the general nature of reasoning with such constructs. In particular, an important underlying theme of the thesis is the general problem of *situated reasoning*—i.e. developing a framework for reasoning in which not all information relevant to the process needs to be made explicit.

My original attraction to situation theory (and channel theory) centred around its possible use in modelling situated representation and reasoning, particularly in relation to recent work reported in the Artificial Intelligence (AI) literature. Many of the central concerns of situation and channel theories are closely related to issues that have recently received attention in this literature, such as indexicality, different perspectives for different agents, focus on a limited part of the world, and interaction between an agent and its environment. In this Introduction, I briefly discuss some of these issues from the perspective of the AI literature before informally outlining some of the properties of situation and channel theories, particularly as they relate to the problem of situated reasoning.

1.1 Formal Models for AI

There is a tradition in AI of using formal tools to model the reasoning of agents and processes.[1] Not surprisingly, formal tools from the philosophy of language have proved extremely useful and popular—many of the concerns of cognition and use of language clearly overlap. While many of the formal frameworks used in the AI literature are based on classical logic, sophisticated techniques have been developed to deal with issues of particular importance (e.g. variants of Hintikka's epistemic logic—e.g. (Halpern and Moses 1985; Levesque 1984)), systems for defeasible reasoning (e.g. (Reiter 1980; McCarthy 1980)), logics for planning and reasoning about action (e.g. (Lifschitz 1987; Georgeff 1987a)). Many of these techniques have found their way back into the philosophy and linguistic communities, leading to new research in these areas.

In AI, it is often the hope that formal models will lead more directly to computational interpretations of the processes involved. However, while the final aim of AI is to build machines that can operate as cognitive agents, the use of formal models to aid in the understanding of the processes involved in cognition is an important end in itself. This is especially true in Cognitive Science, where it is important for any sort of formal model to rest on sound philosophical foundations. Even if such models do not specifically address computational issues, they are still of immense value to the AI community if they address properties of cognition that are of particular concern to AI.

Recent work in the philosophy of language has shown standard truth-conditional logic to be poorly suited to formal semantics. The *dynamic* aspect of language has gained prominence and formal logics designed to address this property have emerged (e.g. (Heim 1982; Kamp 1981; Groenendijk and Stokhof 1991)). What these logics specifically address is the ability of a linguistic utterance to change the *information state* of the hearer—what is important is not under what conditions the propositional content of an utterance is true, but what *information* is provided by that utterance, given the context in which it is made. Similar issues are central to other work in formal philosophy which is of direct interest to AI (e.g. Gärdenfors et al.'s work on belief revision (Gardenfors 1988) and Veltman's (1993) work on defaults and epistemic modalities).[2] Dealing with the dynamics of a changing environment is just one aspect of

[1] The paper by McCarthy and Hayes (1969), where they draw attention to work done in formal philosophy to the AI community, is seminal in this regard.

[2] Work by logicians has also taken a more "informational" slant—e.g. workers in Relevant Logic

current AI research into autonomous agents that must be addressed by any satisfactory formal framework.

Many issues very closely related to those above have recently become the centre of focus in AI research, particularly in the areas of *situated reasoning* and *multi-agent systems.* The range of issues correlates, not surprisingly, with the shortcomings of classical logic when applied to natural language semantics. These include the following:

- The importance of an agent interacting with its environment—the unsuitability of classical formal systems for this task is discussed at some length by Smith (1992);
- The limited scope of perception and reasoning—i.e. agents tend to be concerned only with partial descriptions of their embedding situations (e.g. (Georgeff et al. 1993; Lifschitz 1990));
- The link between mental states and embedding environment: the content of an agent's internal states and the effects of its actions cannot be separated from the way in which it is situated in its environment (e.g. (Rosenschein 1985; Israel et al. 1993));
- The indexical nature of reasoning and the importance of allowing different "perspectives" on the world, particularly in multi-agent models (e.g. (Georgeff et al. 1993; Lesperance and Levesque 1995)).

These are just some of the important issues that a theory of situated reasoning and acting must address. As is described below, these are also central concerns of situation and channel theories. In particular, properties related to the above are not only easily *modelled* but are core properties of the frameworks themselves. While the main content of this thesis does not specifically address the above issues from the perspective of AI agents, the nature of situated reasoning—and the way in which situation and channel theories cater for its central concerns—is an important theme throughout. In Chapter 5, I outline a methodology for situated reasoning based on the particular framework developed in this thesis, and briefly return to the issue of using situation and channel theory as potential frameworks for modelling situated agents in Chapter 6.

(Anderson and Belnap 1975; Dunn 1984; Read 1988) and Linear Logic (Girard 1987) claim that these logics are more concerned with information rather than truth.

1.2 Situation Theory

Much of the formal toolbox I use in this thesis has its roots in situation theory, which is a development of the semantic theory originally developed by Barwise and Perry (1983).[3] Situation theory is concerned with issues which have recently been gathering prominence in the cognitive science and AI communities. *B&P* discuss how their theory of situations arose from what they perceived as an inadequacy with first-order logic to deal with the semantics of natural language. More recently, the aims of the theory have broadened in scope, to an attempt to provide a mathematical theory of information (e.g. (Barwise 1989b; Devlin 1991)). Manipulation of information is at the heart of inference and reasoning, and as such, many of the concerns raised by *B&P* are also of central importance to cognitive science and AI.

My own interest in situation theory arose through the possibility of its use as a framework for building formal models of reasoning, as is popular in AI. Many of the formal techniques that appear in the AI literature are based on conventional logic, and many of *B&P*'s concerns for logic as a basis for the semantics of language hold for the sort of tasks that logic is used for in modelling AI processes. In this thesis, I do not explicitly address AI problems except very briefly towards the end. However, the important issues of the models defined in Chapters 3 to 5—particularly the contextual nature of information and information flow—are themes throughout the work described here, and are closely related to the problems addressed by the AI community under the guise of situated reasoning.

1.2.1 Situations

Situation theory is, not surprisingly, about *situations.* A situation is a part of the world—it may be a visual scene, an event, a whole sequence of events, an utterance, a mental state. *B&P* take a very realistic stance toward situations—the ones that matter are very much a part of the world, existing externally to any agent (although the information supported by a situation is heavily dependent on the way a given agent individuates it—see below). However, such issues are not particularly important to the concerns of this thesis. Unlike a possible-world, a situation does not describe the way the world *could* be, but the way the world *is,* or at least the way it is perceived to be.

[3]Throughout the thesis, I will tend to abbreviate "Barwise and Perry" by "*B&P*".

An important property of situations is that they are *partial,* i.e. they do not resolve all issues there are. This is an important way in which situations differ from possible-worlds and partiality plays a significant role in the use of situations within *B&P*'s semantic theory. Other authors have also stressed the need for partial objects to be incorporated into semantic theories (although not always as situations—e.g. see Landman (1986)). Situations (actually, partial possible-worlds) have also been used to alleviate problems inherent in Hintikka's (1962) modal epistemic logic—i.e. the so-called *logical omniscience* problem—whereby an agent is forced to know/believe all logical consequents of its knowledge/beliefs. Levesque (1984) and Muskens (1989) define modal logics of beliefs based on Hintikka's work, but which avoid some of the unwanted patterns of reasoning associated with logical omniscience.[4]

Another important property of situations is that they are first-class intensional objects—i.e. situations have the same status as other objects and individuals—which means that they can stand in relation to objects, be abstracted as part of a complex object, or participate in the theory in other "interesting" ways.[5] Perhaps most importantly, situations support *infons*, which form the most basic "informational unit". An infon consists of a relation, a sequence of arguments and a *polarity*—this last object indicates whether or not the relation holds of the given arguments, according to the given infon. A pair consisting of a situation and an infon forms a proposition—i.e. a truth-bearing object.[6] These objects are described in more detail in the following chapter.

In some respects, situation theory is reminiscent of the *situation calculus* of McCarthy and Hayes (1969), a formalism popular in the AI literature on reasoning about action and change. The situation calculus is a sorted first-order theory, in which one sort is comprised of *situations.* Basically, each predicate in the situation calculus is accorded an extra argument, which denotes the situation in which the appropriate proposition holds. As such, truth in the situation calculus bears some resemblance to the way truth is defined in situation theory, at least when a situation in the latter is taken to be a (partial) state of the world. Hence, we would expect to take formal models of actions and events (as they are formulated in the situation calculus) and easily recast them in a

[4]Note that the logical omniscience problem does not go away by making possible-worlds partial—all that really happens is that a much weaker notion of consequence is involved.

[5]One consequence of the fact that situations can stand in relations is that if we are to model situations set-theoretically, then the appropriate set-theory must violate the Axiom of Foundation—i.e. we require Aczel's (1988) *non-well-founded set theory*. This issue and the occurrence of non-well-founded situations in the actual world is discussed at length by Barwise (1989d).

[6]I do not make much of a distinction between *infons* and *situation-types* in what follows.

situation theoretic framework (especially since situation theory is a much richer formalism, offering many powerful tools for reasoning with situations). Further to this line of thought is the fact that some recent authors, such as Lifschitz (1990) and Georgeff et al. (1993), have argued that some long-standing problems with the situation calculus can be solved by making the situations of the situation calculus limited in scope, effectively making them partial entities. Georgeff et al. are particularly interested in reasoning about events within a multiple-agent environment, characterising the concept whereby an agent is not aware of all that is going on around it.

1.2.2 Individuation and Situatedness

From the outset, situation theory has been concerned with a relativisation of ontology—the objects, properties and relations of a fragment of situation theory are assumed to be with respect to some *scheme of individuation,*[7] which is generally associated with the view of some particular agent. These objects, properties and relations are discriminated by an agent as uniformities across its environment (Barwise and Perry 1983).[8] Any facts that a situation supports must be constructed from objects, properties and relations, and thereby are only meaningful with respect to a particular scheme. Hence, any information represented in the theory is relativistic, in the sense that it is relative to some such scheme.

An even more fine-grained relativisation is obtained by relativising information *within* a particular scheme, i.e. by taking a particular viewpoint into account. This is modelled by Seligman's (1990) notion of a *perspective.* A perspective basically consists of a collection of situations and the information supported by them. However, the same situation may support different information in different perspectives P_1 and P_2, even when P_1 and P_2 are based on the same scheme of individuation. For example, if s is a situation containing a table supporting a vase and a book, the information supported by s as to whether the vase is behind the book, or vice versa, depends on the particular perspective or viewpoint. This sort of relativisation is one aspect of *situated* information—the way an agent is embedded in its environment plays a part in

[7]Devlin (1991) defines a scheme of individuation as follows: "... a way of carving up the world into various "uniformities" that form the basis of our study: individuals, relations, spatial and temporal locations, and further entities ..." (pg. 26).

[8]This notion of individuation as uniformity across situations is formally modelled rather comprehensively by Seligman (1990).

determining what information that agent extracts from the environment.

Another aspect of situated agents is that the content of intensional states, such as knowledge and belief, is dependent on the environmental circumstances in which the agent is embedded. Again, situation theory takes this as a central thesis: the importance of the embedding situations in representation and inference is considered fundamental. This is demonstrated by what Barwise calls the *Situatedness of Content Principle (SCP),* whereby "the semantic content of a representation in general depends on its embedding circumstances. From the *SCP* it follows that valid inference also depends on embedding circumstances." (Barwise 1989e, p. 157).

Finally, a situated agent has much to gain by making use of *environmental regularities.* The environment surrounding an agent satisfies various informational relations; whether these be based on laws of physics or other relations, perhaps arising from social convention, is not of particular concern. For example, the environmental relationship between smoke and fire, in that the presence of smoke indicates a nearby fire, means that a sprinkler-system for combating fire can be attached to a smoke-detector—when the smoke-detector is activated by the presence of smoke, so too is the sprinkler, thereby dousing the offending fire. Of course, it is possible that a particular environmental relationship may be less than robust in certain circumstances—e.g. we would not want to connect a smoke-detector in a crowded Edinburgh pub directly to such a sprinkler system! The possibility of exceptions in environmental regularities is a primary concern of *Channel Theory* (discussed in Section 1.4.3), which provides the central technical framework of this thesis.

The aspects of relativism and situatedness touched on in this section have relevance to much important work currently being performed in the AI community. An extremely active area of AI research concerns the design of autonomous agents, and "situatedness" is playing an increasingly important role. Several prominent programs of research demonstrate this. Agre and Chapman describe a program that "exploits regularities in its interaction with the world to engage in complex, apparently planful activity without requiring explicit models of the world" (Agre and Chapman 1987). Brooks (1986a, 1986b) describes the design and implementation of robots that have minimal internal representation and basically use the external world to represent their own internal state. Rosenschein and Kaelbling (Rosenschein 1985; Rosenschein and Kaelbling 1986) describe *situated automata,* simple machines that can be given a formal epistemic analysis by exploiting the way in which they are connected to the world. Other authors, writing from outside of situation theory, have also argued for the need to take greater account of an agent's situating environment: e.g. Smith (1987) describes the need for an agent's *participation* in the world, and Suchman (1987) argues that instructions for the behaviour of an intelligent agent cannot be separated from its environment—the interpretation of such instructions and the resulting behaviour is critically dependent on the context in which they occur.

Another important area of current AI research involves the interaction and cooperation of different agents. Different agents can, of course, use different schemes of individuation and take different perspectives on their shared environment. The relativistic nature of situation theory, and Seligman's theory of perspectives in particular,[9] seems ideally suited to formally modelling this aspect of agent interaction.[10] This issue, as well as the possibility of using situation theoretic tools to formally model other aspects of situated AI agents, is discussed in greater detail in Chapter 6.

1.2.3 Information and Information Flow

An important slogan for situation theory is that it is a mathematical theory of *information.*[11] Exactly what constitutes information is not really specified by the theory, as this issue is not fully resolved.[12] According to Israel and Perry (1990), the bearers of information are *facts.* Again, what constitutes a fact is not clear, but facts certainly aren't simply propositions—they have certain properties, such as veridicality and causal properties (e.g. (Barwise 1989b, p. 227)), not possessed by propositions (sometimes not even true ones). However, in Israel and Perry's theory of information, the content of a fact is a true proposition. For all intents and purposes to this discussion, I will take an item of information to be the fact that a situation s is of a particular *type.* Such types are known as *situation-types,* and they classify a situation as being a certain way. Situation-types are discussed in greater detail later.

[9] Seligman's theory of perspectives is at the heart of *Classification Theory*, an integral part of channel theory. This is described in Section 1.4.3 below.

[10] Along similar lines, Healey and Vogel (1994) use channel theory to define a situation theoretic model of dialogue that involves these very properties of agent interaction.

[11] Perhaps this is being both overgenerous and unfair—the claim is usually that situation theory is a step *towards* a mathematical theory of information.

[12] The nature of information is discussed by Barwise (1989c), Devlin (1991), and Israel and Perry (1990), amongst others.

Because of the way in which the world is structured, certain facts give rise to other facts. For example, the fact that it is raining outside gives rise to the information that I will be wet when I arrive home after walking from work. As a theory of information, situation theory's main concern is modelling how an agent uses informational relations about its environment so as to function in that environment—given certain information, what other useful information can be inferred about the world.[13] In *Situation Semantics* (Barwise and Perry 1983), the relevant informational relations are those that hold between types of utterances and the types of situations described by those utterances. In a theory of content based on natural regularity, the relevant relations are those that hold between types of intensional mental states and propositions regarding the world. In a theory of situated inference and cognition, such relations relate types of situations, such as that linking smoke-filled situation-types to *on-fire* situation-types.

In situation theory, informational relations of the type described above are called *constraints.* Constraints were introduced by *B&P* and have played a central role in the theory ever since. Constraints are said to license *information flow:* the information supported, or contained, in one situation s "flows" into another situation s'. The information regarding s can be extracted by an agent because of certain information supported by s'. For example, if s is a situation containing a house on fire, and s' is the situation comprising the smoke-filled sky above s, then an agent, on extracting the information that s' contains billows of smoke, can infer that s contains an object that is on fire. In this case, we say that the fact that s' is smokey *carries* the information that s contains fire.[14]

Constraints such as these can be seen as underpinning any inference, including *analytical* (or "logical") relations. At this point, I simply want constraints to be seen as semantic objects which model the structural regularity of the world—i.e. the way in which information of one type ensures the presence of information of another type. The role of constraints in inference is discussed in more detail in Chapter 2, after I introduce channel theory.

[13] That such information is both situated and relative to some scheme of individuation has already been discussed.

[14] Israel and Perry (1990) are specific on this matter: it is *facts* that carry information, and they do so via constraints.

1.3 Information Flow via Regularities

Before discussing specific models of information flow, I want to briefly consider the general concept of *regularity.* Regularities underlie the relationships between situation-types that give rise to information flow. Barwise and Seligman's[15] *channel theory* is an attempt to model the *structure* of regularity, thereby explaining how regularities can be both reliable yet fallible.

My main reason for discussing here some of *B&S*'s ideas on what constitutes a valid regularity is that in later chapters I use channel theory as a model for conditionals and generics. Without at least sketching out the boundaries as to what does and does not count as a regularity, there is no way to judge the breadth of applicability of the model. One of the attractive features of channel theory is that it does not attempt to *reduce* the notion of a regularity to a function of more primitive objects in the ontology, such as possible worlds—regularities are taken to be irreducible entities in themselves. However, regularities do have structure, as will be described in Section 1.4.3.

1.3.1 The Nature of Regularity

Although the concept of regularity is central to the foundations of channel theory, the theory itself makes very weak claims regarding it. The first, and main, claim that the theory makes is: Regularities exist. That is, regularities between objects and their types exist in the world, and are objects that can be discriminated by an agent and used in cognition and inference to make useful predictions about the behaviour of its environment.

The notion of regularity has long played an explanatory part in philosophy and the claim that such things exist is not a particularly strong one. The problematic issue regarding regularities has generally revolved around the attempt to reduce them to more primitive concepts (e.g. using relations between possible worlds); such attempts have generally been deemed unsatisfactory, either including as regularities relationships that do not intuitively fit the description, or ruling out relationships that seem to be acceptable as regularities. Channel theory, however, makes no such reductionist claims: there is no attempt to reduce the concept of regularity to a more fundamental one involving

[15]Henceforth, I will tend to abbreviate "Barwise and Seligman" as "*B&S*".

concepts that are somehow more primitive. A regularity is seen as an intensional object, discriminated by some agent, which holds in the world without requiring an explanation of its behaviour in terms of its components.

This is a great attraction of channel theory: no commitment is required with regard to how regularities arise, yet the theory provides tools that allow rich explanatory accounts of processes that arise in many diverse domains, such as cognitive science, logic and computer science. The main contributions of this thesis involve showing how channel theory can be used to model various formal systems of conditional reasoning that have traditionally relied upon some reductionist commitment of the conditional relationships central to the system in question. For example, systems of conditional logic (e.g. (Nute 1980)) have traditionally made use of possible worlds and "nearest world" selection functions. The channel theoretic logic of conditionals presented in Chapter 3 takes conditional regularities to be primitive, yet the system obtained displays many of the necessary properties of an adequate logic of conditionals.

It may well be hoped that the primitive-ness of regularities within the channel theoretic framework avoids many of the problems that arise in reductionist treatments of conditional reasoning. This is especially true of conditional and generic sentences, where it is well known that the conditional operators involved are highly intensional and semantics for these that rely on reducing the operator in some extensional way often leads to serious problems (e.g. Morreau (1992a) points out some such problems that circumscription runs into when used as the basis for a semantics of generic sentences). The problems of extensionality, especially for generics, is discussed later in the thesis, and it will be shown that the channel theoretic account avoids these particular problems.

The second major claim of channel theory, with regard to regularities, is that they have structure. The structure that *B&S* claim for regularities in no way involves an attempt to reduce the relationship involved in a regularity to the objects related by that regularity. The claim regarding structure involves a decomposition of a regularity into two aspects: a relationship involving the level of *tokens,* or particulars, and a relationship involving the level of *types* that classify those tokens. This decomposition is made in an attempt to explain the manner in which reliable regularities can still be *fallible,* in that exceptions to a general rule can be observed without invalidating the general rule. Channel theory is an attempt to give an account of regularities that overcomes this very problem.

There are two ways in which channel theory seeks to account for the failure of a regularity. The first involves adding a contextual component to the information-flow model. This allows contradicting regularities to be simultaneously supported by virtue of them being supported by different contexts. The utility of this property is demonstrated in Chapter 4 for the case of generic sentences.

The second, more radical way in which channel theory deals with the possibility of exceptions to general rules is by the decomposition of a regularity into type-level and token-level components. One of the basic tenets of channel theory is that it is not only the relationship between two types that plays a part in the definition of a regularity, but also the connections between tokens that are classified by those types, and the relationship between the two levels. This decomposition allows an account of exceptions to be easily given—a general rule is represented as a type-level relationship, while a particular exception to the rule is given by the absence of the appropriate token-level connection, or by the token-level connection falling outside the domain of influence of the type-level relationship. This notion is made more precise when the formal concepts of the theory are presented.

In some ways, the type-level regularities of channel theory play a very similar role to situation theoretic *constraints*. In the early situation theory literature (e.g. (Barwise and Perry 1983)), constraints involved a relationship between situation-types. However, attempts to provide an account of *conditional constraints*, i.e. constraints that fail in given contexts, were always less than satisfactory. Solutions to this problem generally involved making the type-level relationship itself a situation-type, which was then supported only by certain situations, those which satisfied the context under which the given constraint could be relied upon (e.g. see (Barwise 1986; Nivre 1992; Seligman 1990)). Channel theory can be seen as an attempt to give some flesh to this idea, providing detail about the structure of the particular situations (i.e. channels) that support these constraints. By decomposing regularities into separate token- and type-level relations, *B&S* are able to give an account of when a regularity actually does hold of a given pair of tokens, with respect to the appropriate channel.

1.3.2 What Regularities Aren't

The above discussion has done very little to define what constitutes a regularity. However, there are some relationships which *B&S* most definitely rule out as constituting regularities of the sort they intend. The first case that needs to be ruled out involves chance correlation. A regularity constitutes information about the structure of the world, useful in cognition and prediction, but chance correlation cannot be relied upon in this way. For example, suppose that every time Jerry has visited the UK there has been an election campaign in progress. This does not mean that next time Jerry visits the UK there is any real likelihood of an election campaign being run.[16] Certainly, other information notwithstanding (e.g. being told that an election has been called), Jerry would not plan his next trip to the UK with the strong probability of an election in mind.

It may be tempting to think of regularities as somehow related to causation. However, the mistake of relating regularity and causation *too* directly should not be made. Early work in the analysis of causation attempted to give a reductionist account in terms of regularity. However, this line of thought is well-known to be flawed (e.g. see Lewis (1986)). For example, accounts of causation based on regularity suffer from the *epiphenomena problem,* whereby common effects of a single cause are analysed as related by direct causation. For example, consider a situation where fire causes both smoke and heat.[17] The regularity "Smoke indicates heat" is certainly well supported—any situation in which there is smoke is linked to a situation in which there is fire, which in turn leads to a hot situation—i.e. there is a regularity between smokey situations and hot situations. But the smoke did not cause the heat, nor vice versa. Of course, many regularities *do* arise from direct causation. In fact, it may well be true that *every* instance of direct causation leads to a regularity—the causal link between fire and the smoke it causes leads to the regularity between firey situations and smokey situations.[18]

Even given the problem of epiphenomena (and other problems with the regularity analysis of causation), it may well be that the relationships that we want to be classed as regularity are all somehow related to causation. However, the literature on generic sentences (e.g. (Carlson 1977; Morreau 1992a)) shows that there are many generics that do not arise from causation at all. For example, the *rules and regulations* view of generics argues that many generic sentences are based on conventional rules; one such

[16]Of course, there may be other mitigating factors—Jerry may happen to visit the UK only every four years, and has begun a cycle that corresponds in timing to the running of election campaigns. Of course, even this hypothetical circumstance could not really be relied upon because of the unpredictability of the timing of elections.

[17]To make this regularity more robust, one may want to assume that fire is the *only* cause of smoke.

[18]The notion of "causation" I need for this claim is one between types, rather than between particulars.

generic is the following:

"Bishops move diagonally in chess".

Another class of generic sentences are those based on *descriptive generalisations;* these include sentences such as

"Dogs bark".

We clearly want to consider the relationship underpinning generic sentences such as these as regularities.

It may perhaps be argued that some sense of causation *is* involved in these examples—a chess-piece's bishop-ness "causes" it to be moved diagonally; the fact that d is a dog "causes" d to be an animal that barks. I am not really concerned with how such regularities are viewed since I am not in any way attempting to give a characterisation of the concept of "regularity", but am only trying to outline the concept I have in mind (and also what I understand *B&S* to have in mind) when I use this concept below.

1.3.3 Real and Approximate Regularities

There are a number of important properties of regularities that are specifically addressed by the *B&S* channel theoretic model. As has been described a number of times already, a central concern of *B&S* (1994) is how a regularity can be both *reliable* yet *fallible.* The problematic issue concerns the question, once again, as to what constitutes a regularity—why are some fallible informational relationships (such as "Smoke means fire") classed as regularities while others (such as "Penguins fly") are not ?

The issue proves to be problematic to any attempt to reduce the concept of regularity to more primitive concepts. If the definition is too lax—so as to allow fallible informational relationships to be accepted as regularities—then unacceptable relationships have to be accepted as constituting regularities. On the other hand, taking the view that a fallible relationship is simply an approximation to a "true" regularity—i.e. the limit case being that in which there are no implicit "background assumptions"[19]—has severe problems. Consider again "Smoke means fire". There are numerous situations in which this informational relationship fails—for example, if there is a smoke machine present.

[19] "Background assumptions" are those conditions that are assumed to hold when the constraint is reliable. For example, for the "Smoke means fire" constraint, the background assumptions include the fact that there is not a smoke-machine in the vicinity.

The problem arises because it is generally accepted that—due to it being *impossible* to completely specify all the background assumptions behind such relationships (e.g. see (Suchman 1987; Winograd and Flores 1986))—there is *no* corresponding infallible regularity except the tautological "Smoke and fire means fire", ruling out the possibility of there being any meaningful regularities at all.

The approach to the problem taken by *B&S*, and adopted here, is to accept fallible relationships as *bona fide* regularities and to avoid the temptation to provide a reductionist account of them. The meaningful regularities are exactly those that are discriminated and considered to be meaningful by some agent. This actually fits nicely with situation theory's shift from truth to information—what is of importance is not what makes an informational relationship a regularity, but rather what information can be inferred once a given regularity is accepted. The relativism introduced by the notion of a scheme of individuation permeates the concept of regularities, since the type-level (and token-level) informational relationships discriminated by an agent depend on the types (and tokens) available in the agent's scheme. Hence, different agents will discriminate different regularities in their respective environments.

By rejecting the approach by which a fallible regularity is seen as an "approximation" to some infallible version of it, one is led to a view in which there is no "correct version" of a regularity—each regularity is a regularity *independently* of other regularities that may happen to be related to it. Two regularities may be such that one can be relied upon in a subset of those situations in which the other is reliable—e.g. "Smoke means fire" is less reliable than "Smoke means fire if there are no smoke-machines present"—but the former is no less a regularity than the latter. While the notion of "more reliable" does have a central role to play in the framework for conditional reasoning introduced in Chapter 3, the fact that each regularity can be used independently of any other is crucial in the account of situated reasoning that I outline later in this thesis.

1.4 Models of Information Flow

The important role that information flow plays in situation theoretic analyses of language and cognition was very briefly mentioned above. Barwise and Seligman's *Channel Theory* (Barwise and Seligman 1993; Barwise and Seligman 1994; Seligman 1993; Seligman and Barwise 1993) can be seen as fleshing out the situation theoretic model of information flow, addressing several important issues, such as relativism and context dependence. Most importantly, channel theory is an attempt to give an account of the *structure* of "natural regularities", on which information flow is based, which in turn leads to an explanation of how such relationships can be both reliable and fallible (i.e. can admit exceptions). In particular, channel theory addresses shortcomings in previous models of fallible information flow.

1.4.1 Dretske and Information Flow

The foundational ideas behind situation theoretic constraints come from Dretske's (1981) seminal work on information flow. Dretske defines a probabilistic model of information flow, based on Shannon and Weaver's (1949) quantitative model, that attempts to provide the basis for a naturalistic account of knowledge and belief. The information that an object s is in a certain state or satisfies a particular property can provide information about the state of some other object t. This is known as *carrying information:* the fact that s is in the state it is carries the information that t is in its particular state. Dretske reduces the relation of information carriage to probabilistic measures: s being F carries the information that t is G iff the conditional probability of t being G, given that s is F (and some "background knowledge" k), is 1, while the conditional probability that t is G (given k only) is strictly less than 1. This can be loosely rephrased as: s being F carries the information that t is G iff s being F means that t *must* be G (given k), while it is not the case that t must *always* be G (given k).

Some aspects of the above definition are worth discussing. The first concerns the property that the conditional probability be 1. The value 1 could theoretically be replaced by some "threshold" probability. However, Dretske insists that the "carries information" relation should satisfy what he calls the *Xerox Principle:* if A carries the information that B, and B carries the information that C, then A carries the information that C. Since the conditional probability of C given A is multiplicative of those of C given B and B given A, then the Xerox Principle can only be satisfied if the "threshold" probability is 1. The second point concerns the fact that the probability of t being G (i.e. not conditional on the state of s) must be *less* than one. This means that the relation of carrying information is a stronger one than the material implication of classical logic: in classical logic, $A \supset B$ holds whenever B is valid (or whenever A is unsatisfiable). This is one of the foundational issues of certain logics with *intensional* implication relations, such as *relevant logic* (Anderson and Belnap 1975; Dunn 1984; Read 1988) and *linear logic* (Girard 1987)[20]—such logics are claimed by their authors to be more concerned with information, rather than truth-conditions, and this seems to be supported by this particular property.[21]

Dretske's information-flow theory of content and representation, based on the information-theoretic framework, has been criticized as being unable to adequately explain the possibility of *misrepresentation*—in Dretske's model, the information content of a mental state (representing an external world) becomes too "weak" to ever be incorrect.[22] The problems with Dretske's account stem from an inadequate treatment of exceptions to the information flow relation—given the possibility of exceptions, how can the conditional probability of t being G given that s is F ever be 1? Dretske attempts to circumvent the problem by appealing to *channel conditions,* which amounts to the disclaimer that the above-mentioned conditional probability is 1 *under "normal" conditions.* In some ways, channel theory can be seen as providing technical flesh to this informal notion.

1.4.2 Situation Theoretic Constraints

Before turning to the foundations of channel theory, it is appropriate to briefly discuss situation theoretic constraints. *B&P* take constraints as imposing structure on reality—i.e. constraints are law-like relationships that reside in the world and which cognitive agents can make use of in getting about in the world. These law-like entities exist because "what happens at one place and time must contain information about what has happened or will happen, elsewhere and elsewhen" (Barwise and Perry 1983, p. 94). In the *B&P* model, a constraint is modelled as a relation between two situation-types T and T'—the fact that a situation s is of type T carries information that some other situation s' is of type T'.

Unlike Dretske, *B&P* do not attempt to reduce the law-like dependencies to other, somehow more primitive entities, such as conditional probabilities. The constraints are

[20] Note that not all non-truth-functional implication operators satisfy this property—for example, strict implication $\Box\!\!\rightarrow$ in modal logics satisfies the property that $A \Box\!\!\rightarrow B$ whenever B is valid.

[21] I briefly return to relevant logic and its relationship to the channel theoretic model of information flow in Section A.4.2.

[22] See Fodor's (1990) criticism's, which uses the *disjunction problem* argument. Dretske's proposed fix, distinguishing a "training" stage from a "prediction" stage, is criticised by Koons (1994), who proposes a solution to the problem within Dretske's framework by using a non-standard probability theory.

assumed to hold "out there" in the world, and an important aspect of an intelligent being's interaction with its environment involves discovering and making use of constraints to infer information that it does not necessarily have immediate access to (e.g. by direct perceptive means). Constraints are primitive, though structured, entities. Note that constraints, being situation theoretic entities, inherit important properties from the situation-types they involve. As with situations, the fact that they are objects in the world does not mean that there are not important relativistic considerations. For example, the format of a constraint is dependent on the appropriate scheme of individuation with respect to which the relevant situation-types are defined—changing the scheme can change the constraint. Hence, the world is structured differently for different agents. Of course, this is part and parcel of the situation theoretic approach to modelling cognitive agents.

Of the several sorts of constraint classes that *B&P* identify (see (Barwise and Perry 1983, p. 97)), the most interesting are the *conditional constraints.* Conditional constraints involve *fallible* relationships between situation-types—they may well be localised in their applicability, or contain exceptions under certain "unexpected" conditions. For example, while the informational relationship "Smoke means fire" is in general a reliable one, it is not so in a crowded Edinburgh pub or in the presence of a smoke-machine. Even so, such constraints can still be extremely useful to a cognitive agent. For example, if an agent is inserted in an environment where the "abnormal" conditions do not occur, then the constraint becomes a reliable one—even if the abnormal conditions sometimes arise, if they are infrequent enough then the constraint can still be considered robust enough to be useful.

Various attempts have been made in the situation theoretic literature to represent the conditionality of constraints. *B&P* reduce conditional constraints to unconditional ones by adding the "background assumptions" to the antecedents of constraints. Barwise (1986) requires the informational relation for conditional constraints to be a three-place, rather than two-place, relation—the third argument-place consists of the situation-type under which the constraint (modelled by an informational relation between the other two arguments) is reliable. Nivre (1992) models a conditional constraint as *nested* inside an unconditional constraint—i.e. there is a two-place unconditional informational relation that relates the background conditions to the two-place conditional constraint. Each of these approaches can be seen as an attempt to capture Dretske's notion of "channel conditions". The problem with each approach is that the background conditions must somehow be made *explicit* for the conditional aspect to be effectively captured, which goes against the grain of important recent ideas in cognitive science (e.g. (Suchman 1987; Winograd and Flores 1986)).

Seligman (1990) takes a somewhat different approach—a conditional constraint is just one part of a perspective on the world. A *perspective* models a particular way of viewing the world and the formal concept itself is localised to a specific set of situations and types. As such, the background assumptions of a constraint are not spelt out explicitly, but the fact that they limit the applicability of the constraint is captured by the limited scope of the perspective. This perspectival model of conditional constraints can be seen as a step towards *B&S*'s theory of channels, which accords a central role to Dretske's notion of a communication channel.

1.4.3 Channel Theory

Recently, Barwise and Seligman have been working on a mathematical theory that begins with the intuitions of the information flow models of both Dretske and *B&P*. At the heart of the theory are *channels,* objects based on Dretske's intuitive notion of a communication channel. A channel is structured in a way outlined below, and part of this structure involves objects very much like situation theoretic constraints. However, a channel also encompasses the notion of "normal conditions", reminiscent of Dretske's appeal to such to address the issue of exceptions. This notion is fleshed out somewhat, but does not require any explicit representation of background assumptions, as did previous situation theoretic treatments of conditional constraints. Instead, connections at the level of situations, as well as informational relations at the level of situation-types, are employed in a manner described below. While channel theory does not attempt to reduce regularities to concepts such as conditional probabilities or relations between possibilities, it does attempt to give an account of the internal *structure* of regularity, thereby explaining how informational relations can be reliable yet at the same time admit exceptions.

The issues discussed here are treated very informally—Chapter 2 contains a presentation of the main concepts of channel theory, which constitutes the main framework for the technical results of this thesis.

Tokens, Types and Classifications

Constraints in situation theory involve informational relations between situation-types. In most accounts of conditional constraints, a third situation-type is used to capture the conditions under which the two-place relation can be relied upon. The basis of *B&S*'s theory of channels, however, stems from the observation that it is also the objects at the level of the situations themselves, i.e. at the level of *particulars*, that play a role in the flow of information. As well as the type-level constraints familiar from the standard situation theoretic model, the channel theoretic account of information flow also contains connections between *tokens,* some of which are situations.[23]

For *B&S* (1994), these token-level *connections* are as much a part of the structure of reality as are the type-level constraints.[24] A collection of type-level constraints and a collection of token-level constraints are the two main components of a *channel,* an object that regulates the flow of information. The final component of a channel is a *classification relation,* which indicates which constraints are applicable to which connections. The fact that a connection c is classified by a constraint γ means that information *flows* along c, information of the sort described by γ. When c is "broken" or "abnormal", then either c lies outside the channel or c is not classified in the expected way—again which of these two possibilities occurs depends on the way in which the connections are individuated.[25]

Context Dependency

Context dependency is an important property of information flow, as pointed out by *B&S* (Barwise and Seligman 1993; Seligman 1993). For example, "Swans are white" is a reliable informational relation when residing in Europe, but nowhere near as reliable in Australia. Even so, a cognitive agent can still make use of an informational regularity, so long as the environment in which it operates falls within the context under which the regularity holds. Context in channel theory is modelled by channels themselves having limited scope—an integral part of a channel $\mathcal{C}$ is a collection C of connections,

[23]Actually, what counts as a token and what counts as a type depends on the way an agent chooses to carve up the world. For example, a situation could play the role of a type, rather than a token, in some carving-up of the world. This is made more explicit when channel theory is introduced in greater detail in Chapter 2. To keep matters simple, I will assume for the rest of this discussion that the tokens are situations and the types are situation-types.

[24]Of course, what counts as a token-level connection depends, as usual, on the particular scheme of individuation, as do the connected situations themselves. I return to this point below.

[25]This issue is treated in detail in Section 2.3.4.

which reflects the context of $\mathcal{C}$. Any connection that falls outside C is not applicable to the information flow supported by $\mathcal{C}$. This notion of context is one important way in which the "channel conditions"—i.e. the conditions under which the information flow in question is reliable—are defined. Another way channel conditions are reflected is via the classification relation of $\mathcal{C}$.[26] Most importantly, the channel conditions themselves are not explicitly represented, as they are in the situation theoretic accounts of conditional constraints described above.

A channel supports a particular sort of information flow; i.e. different sorts of information flow are supported by different channels. For example, the fact that the height of mercury in a thermometer carries information about the surrounding temperature is supported by a different channel to that which supports the regularity "smoke means fire". This, of course, is as it should be—the collections of token-level connections that are applicable to such different informational regularities will obviously not overlap greatly (in fact, for these two regularities they are probably disjoint). Hence, the contexts associated with the applicability of different regularities, as well as the channel conditions, can differ wildly, thereby requiring different channels. This means that the world is full of a multitude of all sorts of different channels, supporting all sorts of different information flow, which again is hardly surprising, since, given *B&P*'s original considerations, the world is highly structured. Some of these channels support analytical, or logical, relationships; others support metaphysical ones; while yet others support the more familiar (from our examples!) "smoke means fire" variety. The important point is that each of these channels defines a different context.

Towards a Model of Situated Reasoning

B&S's channel theoretic model of information flow addresses a number of important issues in situated reasoning—e.g. the distributed nature of information, the localisation of facts and the informational links between them, the reliance of regularities on their circumstances—some of which are discussed in the following chapter (see also (Barwise and Seligman 1993)). Perhaps the most important of these issues, for the purposes of this thesis, is the manner in which a fallible regularity can be *used* for inference independently of its relationship to other, more reliable, regularities. The fact that two

[26]The different ways of modelling the failure of the channel conditions seem to account for important properties of generic sentences, as is discussed in Chapter 3.

regularities may be closely related, with one being applicable in a wider context than the other, was briefly discussed in 1.3.3. In the channel theoretic setting of Chapter 3, this leads to an ordering between channels, with one channel being a *subchannel* of another if the regularity corresponding to the first is less widely applicable than the regularity corresponding to the second.

The subchannel relation is used in Chapter 3 to encode the implicit background assumptions of a channel. It turns out that these assumptions need to be explicated in such a way as to ensure that they affect the operations on channels that lead to the logical system for inferring new channels from old—without the background assumptions being taken into account, unacceptable patterns of reasoning for conditionals and generics are obtained. The relation between a channel $\mathcal{C}$ and other channels encodes the background assumptions behind the regularities supported by $\mathcal{C}$. However, while the nature of the subchannel relation is partially determined by the internal structure of $\mathcal{C}$ and its related channels, the background assumptions encoded by the relation are not explicitly represented within $\mathcal{C}$ at all.

This leads to a framework in which a regularity can be *used* for inference completely independently of the background assumptions upon which it is founded, which is a critical property of a satisfactory model of situated reasoning. In the framework involving the subchannel ordering, this translates into the ability to base inference on a channel without needing to refer to that channel's location in the ordering. This property forms the basis of a methodology for situated reasoning which is described in Chapter 5.

1.5 What This Thesis is About

The ideas described in this thesis began from two different points and met in the middle. One starting point involved an investigation into *B&S*'s claim that channel theory could be used as the basis for a semantic analysis of conditional sentences. In particular, I was interested in exploring the viability of using the operations on channels described in Section 2.3.7 as the basis for a channel theoretic system for reasoning about conditionals. At the same time, I was interested in using channel theory for developing a model of *situated planning.* In particular, the channel theoretic model of fallible regularities, coupled to the concept of a subchannel ordering reflecting the notion of a "reliability" ordering on the associated regularities, seemed to suggest an elegant approach to AI's *qualification problem* (e.g. (Georgeff 1987b)), whereby background assumptions behind a plan cannot be easily specified.

On realising that an acceptable logic of conditionals needs to account for background assumptions of a conditional so as to avoid certain unwanted rules of inference, I realised that the framework I was developing for the situated planning task could also be used for the logic of conditionals task. This offered the potential of a uniform framework being used for a number of important related problems that I collectively refer to as *conditional reasoning*, since they all deal with conditional-like constructs: i.e. conditional sentences, generics and default rules,[27] planning operators. The channel theoretic interpretation of each of these constructs offers distinct advantages for each task.

1.5.1 Outline of the Thesis

In Chapter 2, I present the formal concepts which I will be using in this thesis. After a brief informal presentation of the main concepts of situation theory,[28] I present *classification theory,* a relativistic theory of information that constitutes an important component of channel theory. A classification encompasses a collection of tokens and types, together with a relation indicating which tokens are classified by which types. As such, a classification can be seen as representing a particular viewpoint on the world. I then present the main concepts of channel theory, focussing on how different sorts of exceptions to general regularities can be modelled in this framework in different ways. A channel can be seen as a special type of situation, one that supports information of a conditional (i.e. "*if–then*") nature, and the theory involves exploring the internal structure of channels and using this structure to account for exceptions to the regularities supported by them. In particular, channel theory pays close attention to the role of connections at the level of tokens.

Other properties of channels are also presented, the most important (for the purposes of this thesis) being a number of operations on channels. These can be seen as forming the basis of a calculus of channels, or a channel theoretic logic. This logic forms the basis of the channel theoretic systems of conditional and default reasoning defined in

[27] Another motivation for examining generics and default reasoning within this framework stems from recent work which develop logics for default reasoning from a conditional logic basis (e.g. (Morreau 1992a; Boutilier 1992; Lehmann 1989).

[28] A more precise and comprehensive presentation is contained in an appendix. However, the informal presentation will suffice for the purposes of this thesis.

later chapters.

Chapter 3 investigates the viability of using the channel theoretic logic as a basis for a logic of conditionals. A satisfactory logic of conditionals must invalidate certain classically valid patterns of inference, notably, Transitivity, Strengthening of the Antecedent (or Monotonicity) and Contraposition. The channel theoretic analysis of conditional sentences endows the propositional content of a conditional sentence with a demonstrative content—i.e. the object described by the sentence—which is taken to be a channel. The addition of a context (in the form of a channel) to each conditional alleviates the problem of the unwanted patterns of inference but does not remove it altogether—while the channel reflects the assumed background conditions which are in place when a conditional statement is asserted, it does not represent them in a manner that can be exploited by the channel operations so as to invalidate Transitivity, Monotonicity and Contraposition.

A method of encoding background assumptions is proposed within the channel theoretic framework. This involves defining a *subchannel* ordering between channels, which relates two channels whose internal structure satisfies certain simple conditions. This ordering results in a hierarchy being imposed on any given collection of channels. A hierarchy encodes the assumptions of a given channel $\mathcal{C}$ via the way in which $\mathcal{C}$ is related to other channels in the hierarchy. Most importantly, those assumptions are not explicitly represented in $\mathcal{C}$ at all, which is seen as an improvement on a similar approach of Barwise's (1986) to account for background assumptions in a situation theoretic model of conditionals. The channel operations are modified so as to account for a given hierarchy, resulting in a system which does not support the unwanted patterns of reasoning. Further, it is shown that any "reasonable" hierarchy (i.e. one satisfying certain simple constraints) results in a fairly powerful logic of conditionals. This is shown by considering the various axioms and rules of inference discussed by Nute (1980, 1984) in his survey articles.

In Chapter 4, I consider the possibility of obtaining a logic for generic sentences from the channel theoretic logic of conditionals. This was partially motivated by similar work in the AI and philosophy literature within possible-worlds frameworks (e.g. (Morreau 1992a; Boutilier 1992)). By imposing a tacit assumption that a given collection of channels S is closed with respect to background assumptions—i.e. any background assumption to a channel $\mathcal{C} \in S$ is explicitly represented by some other channel in S

supporting a regularity that somehow contradicts the regularity supported by $\mathcal{C}$—I show that a powerful system for reasoning with generics is obtained. In particular, this logic satisfies the Specificity Principle (also known as the Penguin Principle) by which a more specific regularity overides a contradictory regularity that is the composition of two others.

This chapter also contains some discussion of the channel theoretic analysis of generic sentences. As in the analysis of conditionals, each generic is endowed with a context (i.e. a channel). Context is seen to play an important role in the analysis of generics such as the following:

"The Dutch make good farmers and good sailors".

Such examples cause problems for other accounts of generics, particularly those based on the *normative* view of generics (by which a generic is seen as quantifying over all "normal" individuals).

Morreau (1992a) discusses how a useful logic of generics must also be able to be used to draw inferences about the default properties of individuals. In Chapter 5, I define a *maximal normality* condition which allows the system of Chapter 4 to be used to infer information regarding the classification of tokens (i.e. individuals). The normality condition is required since in the channel theoretic model of generics the validity of a generic is independent of the properties of the associated individuals. I show that the default logic thereby obtained effectively defines a *preferential consequence relation*, in the sense of Kraus et al. (1990). While I conjecture that it also defines a *rational* relation, this remains an open problem.

The channel theoretic account suggests an approach to defeasible reasoning that is somewhat different to traditional default logics. The maximal normality condition effectively picks an "appropriate level" in the hierarchy, with respect to the given information, and inference is performed using the channel found at that level. This approach to default reasoning seems to be aligned to the idea of situating reasoning, whereby reasoning proceeds with imprecise regularities (i.e. regularities which can only be relied upon in certain contexts) rather than reasoning with imprecise information. After outlining a general methodology for defeasible reasoning, the approach is illustrated by an application to the qualification problem (one of the original starting points of the work in this thesis, as described above).

Finally, in Chapter 6 I discuss various avenues for future research. One of these involves exploring the methodology for defeasible reasoning described above. Another important topic involves *situated multiagent systems.* As has been discussed at various points in this introduction, the ideas motivating situation theory and channel theory seem to underly important work currently being performed in this area in the AI community. The relativistic nature of information and information flow provided by classification theory seems particularly well suited to the task of modelling multiple interacting agents.

Chapter 2

Channel Theory: Formal Concepts

In this chapter, I present the basics of *Channel Theory*, Barwise and Seligman's model of information flow. This theory provides the formal tools that I use for the rest of the thesis. Most of the concepts and terminology are taken from (Seligman and Barwise 1993), with further input from (Seligman 1993). The foundational and motivational discussion is based on that of (Barwise and Seligman 1994) and (Seligman 1993).

2.1 Elements of Situation Theory

In this section, I present some of the main concepts of situation theory. I will present only a very small portion of the theory and leave many important foundational issues untouched—these are not the primary concern of this thesis. Situation theory is still a developing theory, although many of the original tenets of Barwise and Perry (1983) still hold. Some of the attempts to define large fragments of the formal theory include the following: (Barwise and Perry 1983; Barwise 1989b; Devlin 1991; Nivre 1992; Barwise and Cooper 1991). The last of these forms the basis for the presentation of the formal concepts below.

The presentation of this section is fairly informal, involving a brief discussion of issues that I consider important to the content of the thesis. Appendix A contains a more formal account of basic situation theory, borrowed pretty much directly from

(Barwise and Cooper 1991).[1] However, the concepts described here should adequately address the issues of situation theory required for the purposes of this thesis. One point to be made before commencing is that, following Barwise and Cooper, I assume the theory of situations outlined below to contain set theory, which in particular implies the existence of functions.

2.1.1 Situations, Infons and Propositions

The main objects of the ontology of situation theory are *situations*. A situation may be an event, a sequence of events, a visual scene, a mental state, a classification of objects, or any of a number of other things. In more general terms, a situation is some part of the world that supports facts, where such facts are dependent on some *scheme of individuation*. As mentioned in the previous chapter, the fact that the ontology of situations is dependent on such a scheme has traditionally been at the heart of situation theory, and is considered to be a central aspect of it. Other important issues regarding situations include the fact that they are first-class objects of the theory and that they are partial, in the sense that they do not resolve all possible issues.

Some authors have attempted to formally model situations using more traditional logical tools, such as partial possible-worlds (e.g. (Escriba 1992; Fenstad et al. 1987; Muskens 1989)); another way of modelling partial objects, or partial information states, is via sets of possible worlds (e.g. (Veltman 1993)). While such frameworks have many attractive logical properties, they do not capture the full nature of situations as they are used in situation theory, particularly the fact that situations are first-class objects of the theory.[2] Throughout, I will take situations to be first-class objects, and not defined in terms of partial worlds or by any other similar means.

The second important sort of objects are *infons*, which are the objects which classify situations. An infon consists of a *relation* r, an *assignment* f and a polarity p, where $p \in \{0, 1\}$—such an infon is written $\ll r, f; p \gg$. The relations from which infons are constructed are also taken to be primitive rather than defined set-theoretically, which results in a fine-grained intensionality in the logic of infons.[3] An assignment is a function

[1] As they acknowledge, Barwise and Cooper's definitions owe much to the work of others, notably Aczel and Lunnon (Aczel 1990; Aczel and Lunnon 1991).

[2] These approaches also involve different foundational commitments, such as the ontological commitment to possible worlds.

[3] Such a move is at the heart of *property theory* (e.g. (Bealer 1982; Turner 1987)).

from *roles* (roughly, argument-place-holders) of a relation to objects of the situation theoretic universe. As a convenient abbreviation, I will usually assume that the roles of a relation consist of a sequence of natural numbers, thereby allowing the assignment to be treated simply as if it were an ordered sequence—in this case, an infon involving the assignment $[1 \mapsto a_1, ..., n \mapsto a_n]$ is written as $\ll r, a_1, ..., a_n; p \gg$.[4] For example, the infon $\ll run, john; 1 \gg$ involves the relation run and object $john$; this infon would represent the fact that some individual named John is running. I assume there to be operations on infons, namely infon-conjunction and infon-disjunction. I will say no more on these for the moment—effectively equivalent operations are introduced in Section 2.2.2 in the form of operations on types.[5]

The truth bearing objects of situation theory are *propositions.* One way of forming a proposition is via the binary type $\models$ which holds between situations and infons. The proposition that a particular situation s supports infon σ is written $(s \models \sigma)$ and s is said to *support* σ. Barwise (1989a) calls such a proposition an *Austinian proposition.* For example, $(s \models \ll run, John; 1 \gg)$ is the proposition that makes the claim that the situation given by s supports the information that John is running—this is true exactly if John is running in s. As is usual, I assume there to be a set of operations, namely conjunction, disjunction and negation, over propositions. These operations interact with the type-operations in ways discussed below.

2.1.2 Parameters, Anchors and Restrictions

Parameters are situation theoretic objects that play a similar role to variables in standard first-order logics. Basically, a parameter fills an arguments place without committing to a specific object—e.g. the infon $\ll run, X; 1 \gg$, where X is a parameter, represents the state of affairs where someone is running, without specifying who it is that is doing the running. In general, any situation theoretic object may be replaced by a parameter. Objects containing parameters that are not bound by an abstraction operation (see below) are known as *parametric;* so we have parametric infons, parametric propositions, etc. Given an object o, we assume the existence of a function $par(o)$ that returns the set of parameters in o that are not within the scope of an abstrac-

[4] As a further abbreviation, I will often write the infon $\ll r, a_1, ..., a_n; 1 \gg$ simply as $\ll r, a_1, ..., a_n \gg$, and the infon $\ll r, a_1, ..., a_n; 0 \gg$ simply as $\neg \ll r, a_1, ..., a_n \gg$.

[5] Barwise and Etchemendy (1990) define a model of inference based on an algebra of infons.

tion operation (see below).[6] Note that some properties, such as truth, are not always straightforward for parametric objects.

An *anchor* f is a function that assigns objects to parameters—the domain of f, which is a set of parameters, is written $dom(f)$ and the range is written $rng(f)$. Given an object o and anchor f, the *application* of f to o, written $o[f]$, is the object that results from replacing any parameter $X \in par(o) \cap dom(f)$ by $f(X)$, the value to which X is mapped by f. If $o[f]$ is non-parametric, then f is said to be a *grounding* anchor for o.

Barwise and Cooper (1991) diverge from previous presentations of situation theory in that they allow restrictions to be attached to any object, rather than to parameters only. Given an object o and proposition p, the *restriction* of o to p, written $o \downarrow p$, is the object just like o except that is it restricted so that an anchor can be applied to $o \downarrow p$ only if it meets the conditions in p in addition to any in o. In particular, $o \downarrow p$ is defined if and only if p is not a false proposition, and if p is true then $o \downarrow p = o$. Restriction is a very powerful operation and has important uses in applications of situation theory.

2.1.3 Abstracts and Types

An extremely important and powerful operation that can be applied to a parametric object is *abstraction.* This operation is similar to the abstraction operation of the λ-calculus but involves greater generality. In particular, more than one parameter may be abstracted simultaneously, a feature that makes the indexing of parameters important—for example, in the Aczel–Lunnon (1991) framework, abstraction is performed over a function from a set of roles to a set of parameters, rather than over the parameters themselves. However, as with assignments, I will usually blur this subtlety by assuming that the roles comprise of a sequence of natural numbers, and simply let the ordering of the parameters themselves determine the appropriate role of each parameter.

Abstraction is performed via a binary operation λ, which takes as arguments an indexed set of parameters $[X_1, ..., X_n]$ and a parametric object o. The resulting object, written $\lambda[X_1, ..., X_n]o$,[7] is called an *abstract,* and is itself a first-class citizen in the situation theoretic universe. Dual to the operation of abstraction is that of *application*

[6] Appendix A contains further details on this function.

[7] I will sometimes drop the square brackets, especially when $n = 1$.

of an abstract to an (appropriate) set of objects. Basically, the result of applying an abstract $\lambda[X_1, ..., X_n]o$ to a sequence $[o_1, ..., o_n]$ of objects is the object $o[X_1 \mapsto o_1, ..., X_n \mapsto o_1]$.[8] For example, applying the abstract $\lambda[X]\ll run, X; 1 \gg$ to the object $John$ results in the infon $\ll run, John; 1 \gg$.

A particularly important sort of abstract is the collection of *proposition abstracts*—i.e. objects formed by abstraction over propositions. These objects are known as *types.* For example, the object $\lambda X(X \models \ll run, john; 1 \gg)$ is a *situation-type,* namely, the type of situation in which John is running.[9] Given a type $\lambda[X_1, ..., X_n]o$ and assignment $[o_1, ..., o_n]$, we obtain a proposition via the (partial) binary operation $(. : .)$[10]—the proposition $([o_1, ..., o_n] : \lambda[X_1, ..., X_n]o)$ is true (roughly) iff the proposition that results from applying $\lambda[X_1, ..., X_n]o$ to $[o_1, ..., o_n]$ is true.[11] I will often generalise Barwise's notation and refer to any proposition of the form $(o : t)$ as an Austinian proposition. As can be seen from the notation above, types can in general be n-ary, for any n. However, I will usually be concerned with simple unary types in what is to follow.

In the theory of classifications described below, types are usually presented as if they are unstructured objects—i.e. the internal make-up of a type, via abstraction involving a proposition, is not explicitly shown. This is a matter of convenience, and is also due to the fact that any such internal structure is not relevant to my enterprise—as far as this thesis is concerned, types are first-class objects and they may as well be unstructured.[12] The definition of an adequate semantic model involving such types, however, would generally require greater detail in this matter to be demonstrated. For example, when using situation theory as a semantic framework for the analysis of certain linguistic phenomena, then the internal structure of the (situation theoretic) types used in the analysis is typically a central part of that analysis. However, in this thesis it will generally be the internal structure of channels and regularities that are of importance, and situation-types will generally be treated as unstructured entities.

[8] Of course, both the parameters in the abstract and the objects to which it is applied are assumed to be indexed, in this case by the indices $1, ..., n$.

[9] In fact, all infons, such as $\ll run, john; 1 \gg$, are themselves seen as situation-types, as is formally captured in Appendix A.

[10] This relation can be seen as an instance of the classification relation of the next section.

[11] This is a little sloppy—there are matters such as *appropriateness* to be considered here. This is explained more carefully in Appendix A.

[12] Actually, Barwise and Cooper treat abstracts (and therefore types) as first-class intensional objects, as can be seen in Appendix A.

2.1.4 Constraints

Constraints were introduced by *B&P* as objects that support the flow of information from one situation to another. To *B&P*, a constraint is an infon of the form $\ll involves, T, T'; 1 \gg$, where T and T' are situation-types. Given a situation s of type T, the constraint above indicates the existence of a situation s' of type T'. For example, the informational relation "Smoke means fire" is represented by a constraint of the form $\ll involves, smokey, firey; 1 \gg$—a situation that is of type *smokey* indicates the presence of some other situation that is of type $firey$. In the *B&P* model, constraints underly all information flow, inference and meaning.[13]

B&P present a taxonomy of different sorts of constraints, the most relevant to this thesis being *conditional constraints,* i.e. constraints that can involve exceptions to the general rule. Notable attempts to provide formal models of conditional constraints include (Barwise 1986) and (Seligman 1990). Barwise and Seligman's theory of channels, described in Section 2.3, can be seen as an elaboration of these attempts, particularly that of Seligman. I will be further investigating constraints and the nature of conditional information flow when I present *B&S*'s theory of channels.

2.2 Classification Theory

A theory of *classifications* was introduced by Seligman (1990) and expanded upon by Seligman and Barwise (1993).[14] A classification is an object that supports various propositions, basically that certain particulars, or *tokens*, are of certain types. Classifications form the basis of a relativistic model of information: the tokens of a classification are parts of the world carved out by some agent in question, and the types are properties which can appropriately be used to classify those tokens.

2.2.1 Basic Concepts

A classification collects various particulars in the world and assigns types to them. The particulars are called *tokens:* they may be just about anything occurring in the world,

[13] Nivre (1992) provides a detailed analysis of communication and meaning in a situation theoretic framework, making heavy use of constraints.

[14] Henceforth, I will tend to abbreviate "Seligman and Barwise" by *S&B*.

such as individuals, objects, situations, brain-states, events, or even other classifications. The types that are assigned to them are any properties from any given criteria: objects may be classified according to size or colour; events may be classified according to duration, or time of occurrence; situations may be classified according to situation-type; logical sentences may be classified by first-order models; thermometers may be classified according to height-of-mercury, or by manufacturer.

Definition *A* classification A *is a structure* $\langle tok(A), typ(A), K^+, K^- \rangle$*, consisting of a set of* tokens *tok(A), a set of* types *typ(A), and set valued functions* $K^+, K^- : tok(A) \rightarrow 2^{typ(A)}$.

In simple cases, a classification will be concerned with classification along a particular dimension (e.g. colour, temperature), although this is certainly not necessary. Following *B&S*, classifications are represented diagrammatically as shown in Figure 2.1—the set of tokens is represented by the bottom plane, the set of types by the ellipse above it, while the classification of a specific token by a specific type is represented by the connecting line.

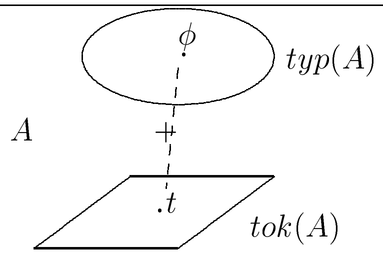

Figure 2.1: A Classification A.

Given a classification $A = \langle tok(A), typ(A), K^+, K^- \rangle$, $a \in tok(A)$ is said to be *positively classified* (resp., *negatively*) by $\phi \in typ(A)$ *in* A if $\phi \in K^+(a)$ (resp., $\phi \in K^-(a)$). Rather than referring to the functions K^+ and K^-, I will usually make use of the following relations $:^+_A, :^-_A$ that can be defined from them:

$a :^+_A \phi$ iff $\phi \in K^+(a)$ in A;
$a :^-_A \phi$ iff $\phi \in K^-(a)$ in A.[15]

These relations are known as *positive* and *negative classification relations,* respectively.

[15] I will tend to drop the subscript from $:^+$ and $:^-$.

As is the case for the functions K^+, K^-, these relations are defined with respect to a given classification A. I will also refer to the positive classification relation as an *of-type* relation—a is said to be *of type* ϕ in A if a is positively classified by ϕ in A.

The way in which the classification relations of a classification have been defined clearly leaves room for both under- and over-specification: i.e. for any token a and type ϕ, it could be the case that *neither* $(a :^+ \phi)$ nor $(a :^- \phi)$ holds, or that *both* $(a :^+ \phi)$ and $(a :^- \phi)$ hold. A classification that satisfies the first of these cases is said to be *partial;* a classification that is not partial is said to be *total.* A classification that satisfies the second of the above cases is said to be *incoherent;* a classification that is not incoherent is said to be *coherent.* A classification that is both total and coherent is said to be *classical.* Throughout the thesis, I will concern myself with coherent classifications.

Classifications form the basis of a relativistic theory of information. The information supported by a classification A is obviously relativised to the tokens and types contained in A, as well as to the classification relations associated with A. This viewpoint avoids any commitment to an objective world outside of the classifying agent, and is most useful when modelling aspects of multi-agent interaction. For example, different classifications can arise over the same collection of tokens and types, which can be used to model different agents having different "perspectives" or viewpoints over the same objects.

2.2.2 Adding Structure to Classifications

The notion of a classification is a simple one and involves very little ontological commitment. In fact, classifications have analogues in model theory, i.e., models—the tokens are the individuals and the types are properties, with the of-type relation being the predication relation. Throughout this thesis, however, I will assume that classifications have some added structure to them, inherent mainly in the types.[16]

Definition *Any set of types typ(A) of a classification A is assumed to be closed under operations of conjunction, disjunction and negation, denoted $\wedge, \vee, \neg$ respectively—i.e. if ϕ, ψ are types in typ(A), then so too are $(\phi \wedge \psi), (\phi \vee \psi), (\neg\phi)$. The type-negation operation is constrained to satisfy the following condition: for any token a and type ϕ, $a :^+ \phi$ iff $a :^- \neg\phi$, and $a :^+ \neg\phi$ iff $a :^- \phi$.*

[16] Seligman and Barwise use the term *perspective* to refer to such classifications—i.e. those whose types have internal structure in the way described here.

The above property of the type-negation operator allows a simplification to the notation since we can do away with the negative classification relation '$:^-$'; this allows us to write the positive classification relation simply as ':'. Hence, given classification $\langle tokA, typ(A), K^+, K^- \rangle$, $\phi \in K^+$ iff $a : \phi$, while $\phi \in K^-$ iff $a : \neg\phi$. Throughout the rest of the thesis, I will adopt this simplification in notation, except where explicitly stated otherwise.

Definition *Any set of types $typ(A)$ is assumed to come equipped with a* type-entailment relation $\leq_A$[17]*—for $\phi, \psi \in typ(A)$, $\phi \leq_A \psi$ means ψ type-entails ϕ. Any such type-entailment relation is assumed to be at least a preorder. The minimum interaction required between the type-operations and the entailment relation of $typ(A)$ is the following: for any $\phi, \psi \in typ(A)$, $\phi \leq_A (\phi \wedge \psi)$ and $(\phi \vee \psi) \leq_A \phi$.*[18] *Given a classification A, the type-entailment relation $\leq_A$ associated with A constrains the classification relation of A as follows: for $a \in tok(A)$ and $\phi, \psi \in typ(A)$, if $a : \psi$ and $\phi \leq_A \psi$ then $a : \phi$.*

The above definition provides the minimum conditions that a type-entailment relation must satisfy. Of course, we may want a particular type-entailment relation to satisfy stronger conditions than these; for example, we may want such a relation to mimic the entailment relation of some logic.[19] In a later section, I will discuss the use of a *logical channel,* which supports the information-flow provided by some logical system, to constrain the entailment relations of certain classifications.

It should be noted that type-entailment can be used to model more than just "logical" entailment between types. For example, the notion of type-entailment can be useful in a multi-agent setting where we want to model different agents having different abilities of discrimination. We may associate a classification A with an agent that cannot distinguish between the colours blue and green—$typ(A)$ would contain the type *blue.or.green*. Another classification B could model the normal case—$typ(B)$ would contain the types *blue* and *green*. A classification that provided a "God's-eye view" of this world would contain all three types, with *blue.or.green* $\leq$ *blue* and *blue.or.green* $\leq$ *green* holding over its collection of types.[20]

---

[17]I will often leave off the subscript when it is clear from the context.

[18]Throughout the thesis, I will tend to assume that type-entailment mimics (zero-degree) *relevant entailment* in the sense of (Anderson and Belnap 1975). Relevant entailment (and its link to channel theory) is discussed in greater detail in Appendix A.4.2.

[19]This involves drawing the obvious analogy between the structured types of a classification and the sentences of some (propositional) language.

[20]Of course, whether one would want to allow a "God's-eye" classification is debatable, but this is another issue.

Finally, work described in a later chapter requires that classifications come equipped with a *type-conflict relation.*

Definition *A* type-conflict relation $\perp_A$ *over a set of types* $typ(A)$ *is an irreflexive, symmetric relation that holds between two types that are mutually incompatible. The minimum conditions required of such a relation are the following:*

$\phi \perp_A \neg\phi$, *for all* $\phi \in typ(A)$*;*
if $\phi \perp_A \psi$ *and* $\psi \leq_A \tau$ *then* $\phi \perp_A \tau$*, for all* $\phi, \psi, \tau \in typ(A)$*;*
if $\phi \perp_A \psi$ *then either* $a : \phi$ *or* $a : \psi$ *fails to hold, for* $a \in tok(A)$*,* $\phi, \psi \in typ(A)$*.*

Once again, as with type-entailment relations, there will be occasions where tighter restrictions on a classification's type-conflict relation will be required. For example, we could require that

$\phi \perp_A \psi$ iff $\neg\phi \leq_A \psi$ and/or $\neg\psi \leq_A \phi$.

2.2.3 Operations on Classifications

So far, I have implicitly taken classifications to be somehow unstructured at the token-level—any logical structure has been present only at the type-level, through the use of type-entailment relations and type-operations. However, this does not allow the representation of complex propositions involving more than one token. For example, let A be a classification containing tokens a, b and types ϕ, ψ; there is no way of representing the fact that A supports the proposition that is the conjunction of $(a : \phi)$ and $(b : \psi)$. Of course, A can support each of the individual propositions, but we would like some way of representing the complex proposition within the classification-theoretic framework.

Seligman and Barwise (1993) define a number of operations on classifications which can be used for this purpose. For example, the following operation of *sequential conjunction* allows the representation of the conjunction of (Austinian) propositions.[21]

Definition *(Seligman and Barwise 1993) Given classifications* A_1 *and* A_2*, the* sequential conjunction $A_1 \otimes A_2$ *is the classification whose tokens are* $\{\langle \wedge, a_1, a_2\rangle \mid a_1 \in tok(A_1), a_2 \in tok(A_2)\}$ *and types are* $\{\langle \wedge, \phi_1, \phi_2\rangle \mid \phi_1 \in typ(A_1), \phi_2 \in typ(A_2)\}$*, and whose classification relation is defined as follows:*

[21] Seligman and Barwise actually define a more general concept allowing the conjunction of an arbitrary collection of classifications.

$\langle \wedge, a_1, a_2 \rangle :^+ \langle \wedge, \phi_1, \phi_2 \rangle$ *in* $A_1 \otimes A_2$ *iff* $a_1 :^+ \phi_1$ *in* A_1 *and* $a_2 :^+ \phi_2$ *in* A_2*;*

$\langle \wedge, a_1, a_2 \rangle :^- \langle \wedge, \phi_1, \phi_2 \rangle$ *in* $A_1 \otimes A_2$ *iff* $a_1 :^- \phi_1$ *in* A_1 *or* $a_2 :^- \phi_2$ *in* A_2*.*

The classification $A \otimes B$ effectively supports (Austinian) propositions of the form $(s : \phi) \wedge (t : \psi)$, where $(s : \phi)$ and $(t : \psi)$ are propositions whose truth is determined in A and B respectively. Seligman and Barwise define operations on classifications corresponding to disjunction, negation, complementation and other useful properties. That said, in this thesis I will tend to reflect all logical structure at the type-level—however, it is important to know that the more general framework is available.[22]

A characterisation that will prove useful in later chapters is the following concept of one classification being a *subclassification* of another. Actually, I define two such concepts, closely related to each other.[23]

Definition *Classification* A *is a* subclassification *of classification* B*, written* $A \sqsubseteq B$*, iff* $tok(A) \subseteq tok(B)$*,* $typ(A) \subseteq typ(B)$*,* $\leq_B$ *agrees with* $\leq_A$ *when restricted to* $typ(A)$*, and for each* $a \in tok(A)$ *and* $\phi \in typ(A)$*, if* $(a : \phi)$ *holds in* A *then* $(a : \phi)$ *holds in* B*. If this last condition is strengthened to a biconditional (i.e.* $(a : \phi)$ *holds in* A *iff* $(a : \phi)$ *holds in* B*), then* A *is said to be a* restriction *of* B*, written* $A \angle B$*.*

Note that A being a restriction of B is a slightly stronger condition than A being a subclassification of B, in that every restriction of a classification is a subclassification of it, but not vice versa. For example, suppose A is a subclassification of B, with $a \in tok(A)$, $\phi \in typ(A)$, $(a : \phi)$ holds in B, but $(a : \phi)$ does not hold in A—then A is not a restriction of B. The rationale for introducing the notion of a restriction (Seligman and Barwise (1993) only define subclassification) is that I want to capture the concept of an agent extending the set of types it uses for classification without altering the way it uses the original set of types. Both the above concepts clearly impose partial orderings on classifications whenever the type-entailment relations are themselves partial orders. I write $A \sqsubset B$ for $A \sqsubseteq B$ and $A \neq B$, where equality between classifications A and B is defined as follows: $A = B$ iff $tok(A) = tok(B)$, $typ(A) = typ(B)$, and $A \angle B$.

One final pair of concepts that will prove useful later on (in the definition of a *channel,* Section 2.3.3) is the following.

[22] Appendix A.2 contains some further illustrations of the use of classification operations to generalise concepts defined below.

[23] Subclassification needs to be extended in the presence of classification operations such as $\otimes$. This is briefly discussed in Appendix A.2.

Definition *A* bi-function $f : A \rightrightarrows B$ *from classification* A *to classification* B *is a pair of functions* $\langle f^{\wedge}, f^{\vee} \rangle$ *with* $f^{\wedge} : typ(A) \to typ(B)$ *and* $f^{\vee} : tok(A) \to tok(B)$.

Definition *A bi-function* $f : A \rightrightarrows B$ *is a* homomorphism *from classification* A *to classification* B *iff:*

- *if* $a :^{+} \phi$ *in* A *then* $f^{\vee}(a) :^{+} f^{\wedge}(\phi)$ *in* B*;*
- *if* $a :^{-} \phi$ *in* A *then* $f^{\vee}(a) :^{-} f^{\wedge}(\phi)$ *in* B*.*

2.2.4 Situation Theory and Classifications

Seligman and Barwise (1993) show how basic situation theory can be recovered from classification theory. Situations are equated with a certain sort of classification, so-called *star-classifications.* In the reconstruction of situation theory, the types of such classifications are relations and the tokens are finite sequences of elements of some underlying set. If A is such a classification, then the situation corresponding to A supports the infon $\ll r; a_1, ..., a_n; p \gg$ (where $p \in \{+, -\}$) if and only if $a_1, ..., a_n :^{p} r$ in A.

In this thesis, I am more concerned with the relationship between situation theory and classification theory by which we have classifications whose tokens are situations and whose types are situation-types. I have already introduced the fragment of situation theory that I am concerned with in this thesis (see Section 2.1). Taking tokens to be situations and types to be situation-types is straightforward, except for the case where parameters are involved. Disallowing parameters would severely curtail the expressive power of the channel theoretic models of later chapters. In particular, parameters are needed if the full expressive power of the *GEN* operator for representing the meaning of generic sentences is to be captured (see Section 4.1.3). However, the treatment of parameters within the channel theoretic framework needs to be handled with some care, especially if we are to model the behaviour of the *GEN* operator, and of discourse referents in DRT. Section A.3 contains one treatment of parameters within the channel theoretic framework that results in the required behaviour.

2.3 Channel Theory

Channel Theory (Barwise 1993; Barwise and Seligman 1993; Barwise and Seligman 1994; Seligman 1990; Seligman 1993; Seligman and Barwise 1993) is a theory of *information-*

flow. The ideas behind channel theory have their roots in situation theoretic *constraints*. The whole concept of the analysis of information-flow had its beginnings in Shannon and Weaver's (1949) quantitative analysis of information-flow for communicating systems, which led to Dretske's (1981) seminal work. Many of the central concepts of channel theory, in particular its explicit focus on the possibility of *error* in information-flow, can be seen as an explicit attempt to address the sort of problems that arise in Dretske's account of a naturalistic theory of content based on information-flow.

2.3.1 Types, Tokens and Information Flow

As I will discuss below, there are several properties of channel theory that contribute towards its attractiveness as a model of information-flow. Perhaps the most crucial aspect, however, is the careful distinction between the token level and the type level. It is this distinction that leads to the theory's ability to account for the fallibility of general regularities, which is a central contribution of the theory.

A channel supports information flow of a certain sort. As will be seen, this is reflected in the fact that a channel has an associated set of tokens and types, those which are appropriate to the regularity in question—this plays the part of a "context" for the information flow. Basically, the scope of a channel defines the "normal conditions" under which a regularity can be relied upon. For example, regularities involving laws of physics generally assume some "ideal" conditions—a channel supporting such a regularity will be defined by these conditions, although such conditions are not explicitly represented in the channel. The information flow supported by a channel is relativised in another important way: as is the case in the theory of classifications outlined above, tokens and types exist only with respect to scheme of individuation. Hence, the existence and form of a regularity, both at the token-level and at the type-level, depends very much on the way the tokens and types are discriminated.

The distinction between tokens and types extends to the model of regularities themselves—it is not just the regularity between a pair of types that plays a role, but also the connection between a pair of tokens that support those types. It is not always clear-cut as to what can play the part of a connection: in an electrical circuit, it may be the conducting wire between two components; when talking about the "connection" between a thermometer and patient, however, the actual connection in question is more abstract. Barwise and Seligman (1994), however, are quite specific about their claims towards the realism of connections: connections between tokens are part of the natural order of things.

It is the connection between a pair of tokens that actually supports a given instance of information flow—figuratively, we could view the information as flowing along the connection, from one token to the other—while the link between types indicates the type of flow that takes place. The direction of the flow is in the opposite direction to inference. For example, consider the usual "Smoke means fire" regularity: the information that a situation s' is firey flows into a smokey situation s, allowing an observer of s to infer that there is a firey situation nearby. This is illustrated in Figure 2.2. Also, it is the *tokens* that are the information bearers—s carries the information regarding s' (i.e. that it is firey).

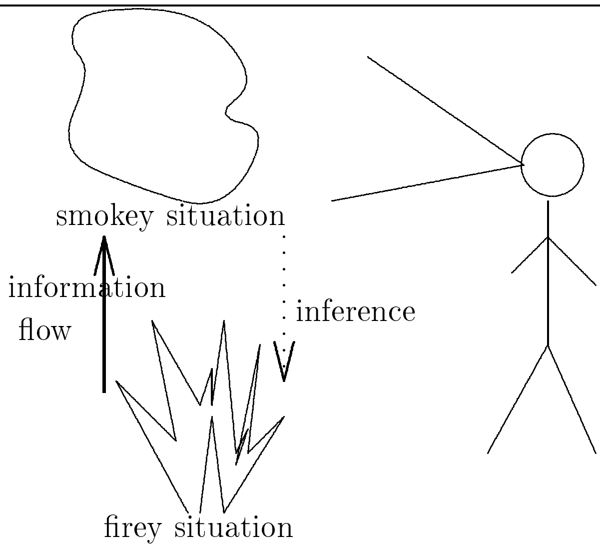

Figure 2.2: Smoke means fire: direction of information flow.

The points made in this section outline some of the ideas behind *B&S*'s channel theory. The rest of this chapter involves introducing the formal concepts of the theory, and expanding on the points made above in the light of these concepts.

2.3.2 Information Links

Section 2.2 introduced *B&S*'s theory of classification. As described there, classification theory can be seen as a relativistic theory of information. Channel theory is concerned with relativistic information *flow,* i.e., how information supported in one classification carries information regarding another. The previous section introduced the notion of connections between tokens and regularities between types, and the manner in which information "flows" along these connections. This notion is formally captured by the notion of an *information link.*

Definition *An* information link *(or just* link*)* $L : A \rightrightarrows B$ *from classification* A *to classification* B *consists of an* indicating *relation* $L^{\wedge}$ *on* $typ(A) \times typ(B)$ *and a* signalling *relation* $L^{\vee}$ *on* $tok(A) \times tok(B)$. *I will sometimes write* $\phi \rightarrow_L \psi$ *for* $\langle \phi, \psi \rangle \in L^{\wedge}$*; when this is the case, we say that* ϕ indicates ψ in L. *Similarly, I will sometimes write* $a \overset{L}{\mapsto} b$ *rather than* $\langle a, b \rangle \in L^{\vee}$*; when this is the case, we say that* a signals b in L.[24]

An information link is a model of the flow of information between two classifications. The signalling relation models the connections between tokens, while the indicating relation models the type-level relations. The information flow modelled by a link is relativistic in that it is dependent on the tokens and types discriminated in the linked classifications, as well as the particular signalling and indicating relations themselves. In general, an information link is to be seen as modelling information flow of a specific sort—every tuple in the indicating relation is applicable to every element of the signalling relation, so the information-flow modelled by a particular link is correspondingly constrained. (Note that the direction of the relations is in the direction of *inference* rather than *flow* as indicated by Figure 2.2.)

As a simple illustration, consider the usual "Smoke means fire" example. The fact that a situation s is "smokey" carries the information that some other situation s' is "firey". This is modelled as follows. There is some classification S that classifies situations according to whether or not they are smokey: i.e. $tok(S)$ consists of a set of situations, and $typ(S)$ contains the single type *smokey.* Another classification F classifies situations according to whether or not they are firey: $tok(F)$ also consists of a set of situations and $typ(F)$ contains the single type *firey.* The regularity "Smoke means fire" is modelled as a link L between S and F. A situation s that is smokey is linked to a particular situation s'—s' is the firey situation that causes s to be smokey; i.e. s' being firey causes s to be smokey. In the link L, this is represented by the pair $\langle s, s' \rangle$ in the signalling relation of L and the pair $\langle smokey, firey \rangle$ in the indicating relation. This models the flow of information from s' to s—the fact that $(s : smokey)$ holds in S

[24] I will usually leave off any reference to L in the above notation when it is clear from the context.

carries the information that $(s' : firey)$ holds in F.

The important motivation behind channel theory is the ability to model error in information flow. The token/type distinction, which is present in the concept of information link, allows the following characterisation of unexpected behaviour in information flow.

Definition *Let* $L : A \rightrightarrows B$ *be a link. A* pseudo-signal *for a constraint* $\phi \rightarrow_L \psi$ *in* L *is a token* $a \in tok(A)$ *such that* $(a : \phi)$ *holds and there is no* $b \in tok(B)$ *such that* $a \overset{L}{\mapsto} b$ *in* L.[25]

A pseudo-signal can lead to error if an agent is not aware of its existence. For example, consider again the link L above modelling the regularity between smokey and firey situations, and suppose that s is a smokey situation that is caused by a smoke machine. Within the signalling relation of L, s is not linked to any situation—i.e. there is no situation s' whose firey-ness causes s to be smokey. In this case, s is a pseudo-signal in L. However, if I am in the vicinity of s and notice its smokey-ness, without realising that this state was caused by a smoke machine, I may well come to assume that there *is* some nearby firey situation and call the fire-brigade. It is in this way that pseudo-signals can lead to error.

Pseudo-signals, and the manner in which they lead to error, are discussed in greater detail in Section 2.3.4, where they are compared to *exceptions*. This latter concept constitutes a more interesting characterisation of error, which becomes available when we add classification relations to links, leading to the concept of a *channel* (Section 2.3.3).[26] A difference between a pseudo-signal and an exception is that the latter can lead to error even when there is a connection between a signal token and a target. This circumstance arises when the information carried by the signal a about some target b disagrees with the information directly supported by b in the appropriate classification.

Definition *A link* $L : A \rightrightarrows B$ *is* sound *if, for all* $a \in tok(A), b \in tok(B)$ *such that* $a \overset{L}{\mapsto} b$ *and all* $\phi \in typ(A), \psi \in typ(B)$ *such that* $\phi \rightarrow_L \psi$*, if* $(a : \phi)$ *holds in* A *then* $(b : \psi)$ *holds in* B.

The fact that a link is sound or unsound is not reflected within the internal structure

[25] *S&B* also define a dual concept, that of *pseudo-target*. The reason for my omitting the dual concept is discussed below.

[26] Actually, there is a bit more to the difference between a channel and a link than this classification relation. These differences are discussed later in this chapter. Note that a link in the sense of *S&B* (1993) is more or less the same as the concept of a channel as it first appeared in the situation theoretic literature (e.g. (Barwise 1993; Barwise and Moss 1991)).

of the link itself—the succedent classification needs to be checked before any unsoundness shows itself. However, the introduction of a classification relation, which leads to the concept of a channel, provides for an extra degree of freedom which allows soundness to be be modelled more directly, as shown in the next section.

2.3.3 Channels

In this section, I present *S&B*'s concept of a channel. Channels are objects that support regularities—underpinning any regularity is a channel of some sort. Channels are related to the concept of information link introduced above but are seen as primary to it—any information link is obtained from a channel by removing some added structure that is available in the channel. As is the case with information links, a channel separates regularities into distinct token- and type-levels, but also involves a classification relation between the objects at these two levels. Further, channels have finer-grained intensionality than links—connections between tokens and types are primitive objects and are not identified with the tokens/types they connect. The definition of a channel involves the notion of a homomorphism between classifications, which has consequences that are discussed below.

Definition *Let A and B be classifications. A* channel *from A to B is a triple*

$$\langle left(\mathcal{C}), \mathcal{C}, right(\mathcal{C}) \rangle$$

where $\mathcal{C}$ is a classification, and $left(\mathcal{C}) : \mathcal{C} \rightrightarrows A$ and $right(\mathcal{C}) : \mathcal{C} \rightrightarrows B$ are homomorphisms from $\mathcal{C}$ to A and B respectively.

The types of $\mathcal{C}$ are called constraints *and the tokens are called* connections*. Given a constraint γ, $left(\mathcal{C})^\wedge(\gamma)$ is called the* antecedent of γ *and is sometimes denoted by $ante_{\mathcal{C}}(\gamma)$, and $right(\mathcal{C})^\wedge(\gamma)$ is called the* succedent of γ *and is sometimes denoted by $succ_{\mathcal{C}}(\gamma)$. Given a connection c, $left(\mathcal{C})^\vee(c)$ is called the* source of c *and is sometimes denoted by $source_{\mathcal{C}}(c)$, while $right(\mathcal{C})^\vee(c)$ is called the* target of c *and is sometimes denoted by $target_{\mathcal{C}}(c)$.*[27]

Figure 2.3 illustrates diagrammatically the concept defined above. As an example of a channel, consider the *Rey* channel $\mathcal{R} : T \Rightarrow P$ which relates a thermometer's reading to a given patient's body-temperature.[28] The classification T classifies thermometers

[27] The subscripts are dropped when the channel in question can be determined from the context.

[28] This example is taken from (Barwise and Seligman 1994).

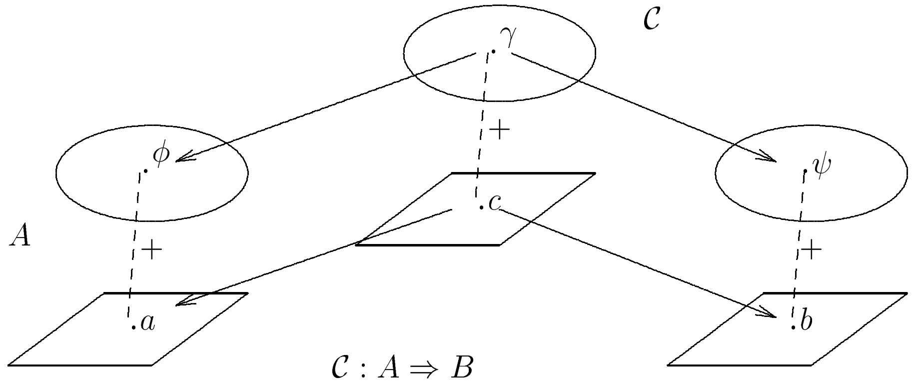

Figure 2.3: A Channel $\mathcal{C} : A \Rightarrow B$.

by height of mercury (in centimetres) and the classification P classifies patients by body-temperature (in degrees Celsius). The *Rey* channel is illustrated in Figure 2.4. The connections of $\mathcal{R}$ are between thermometer-tokens and patient-tokens—a particular thermometer-token t is related to a particular patient-token p if t indicates something about p (rather than some other patient!). The constraints of $\mathcal{R}$ relate heights of mercury (which classify thermometer-tokens) and temperatures (which classify patients)—a particular height-type *height* is related to a particular temperature-type *temp* if (in general) a reading of *height* indicates a temperature of *temp*.[29] The fact that a particular constraint γ (i.e. height-temperature link) is related to a particular connection c (i.e. thermometer-patient connection) is given by γ classifying c in the classification $\mathcal{R}$. When this is the case, then (by the Principle of Harmony described below) the fact that the thermometer t (where t is the source of c, from the diagram) contains mercury at height *height carries the information* that the patient p has a body-temperature of *temp* degrees.[30]

The fact that the definition of channel makes use of homomorphisms has several

[29] This should be qualified slightly: a height of *height* indicates a temperature of *temp* under the "normal conditions" reflected by $\mathcal{R}$. The fact that a channel reflects the "normal conditions" of a regularity (and thereby provides a characterisation of error) assumes prominence below.

[30] In general, given a connection c' in $tok(\mathcal{R})$ and constraint γ' in $typ(\mathcal{R})$, $\mathcal{R}$ can be seen as supporting a conditional of the following form:

> "If $source(c')$ contains mercury at height $ante(\gamma')$ then $target(c')$ has body-temperature $succ(\gamma')$."

This is the case even if c' is not classified by γ' in $\mathcal{R}$—in this case, the antecedent (and consequent) of the conditional sentence may not hold (i.e. the sentence may be a counterfactual). This is the basis of the channel theoretic analysis of conditional sentences described in the following chapter.

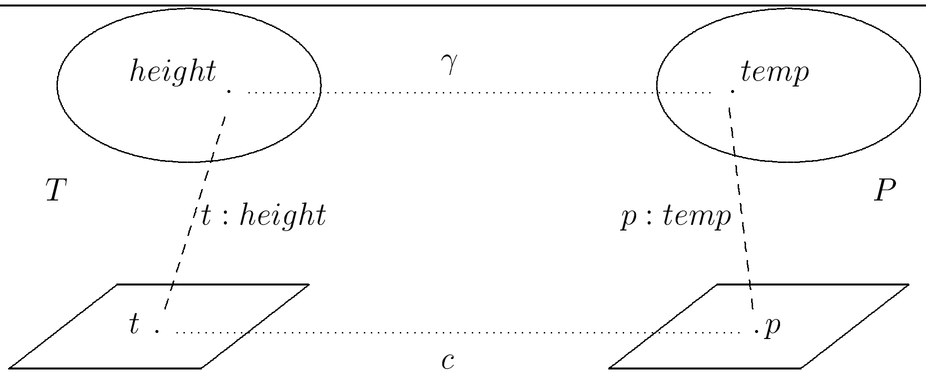

Figure 2.4: Illustration of the *Rey* channel $\mathcal{R} : T \Rightarrow P$.

consequences. The most important is that a connection's being classified by a constraint means that information-flow occurs: i.e. for a channel $\mathcal{C} : A \Rightarrow B$, if $(c : \gamma)$[31] holds in $\mathcal{C}$, then it *must* be the case that $(left(\mathcal{C})^{\vee}(c) : left(\mathcal{C})^{\wedge}(\gamma))$ and $(right(\mathcal{C})^{\vee}(c) : right(\mathcal{C})^{\wedge}(\gamma))$ both hold in their respective classifications. The idea is that (positive) classification of a connection by a constraint models the fact that information-flow *has occurred*—i.e. the source *is* of the type given by the antecedent of the constraint, and (therefore) the target *is* of the type given by the succedent. Seligman (1993) refers to this as the **Principle of Harmony,** which he represents diagrammatically as shown in Figure 2.5.

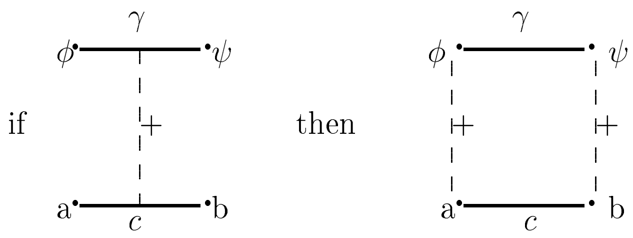

Figure 2.5: Principle of Harmony.

A similar principle exists for negative classification, of course: if $c :^{-} \gamma$ in $\mathcal{C}$, then it

[31]Given the notation so far, positive classification should be distinguished from negative classification for channels. However, I will sometimes leave off the superscript; such cases are assumed to correspond to positive classification. I return to this issue below.

must be the case that $(left(\mathcal{C})^{\vee}(c) :^{-} left(\mathcal{C})^{\wedge}(\gamma))$ and $(right(\mathcal{C})^{\vee}(c) :^{-} right(\mathcal{C})^{\wedge}(\gamma))$ both hold. However, I will not make use of this property. My reason is that, in Section 2.3.7, I introduce a *Contraposition* operation that defines a channel $\mathcal{C}^* : B \Rightarrow A$ for any channel $\mathcal{C} : A \Rightarrow B$. The contraposition channel $\mathcal{C}^*$ satisfies the Principle of Harmony (for positive information flow) exactly when $\mathcal{C}$ satisfies the dual principle (for negative information flow). Hence, I replace the need to talk about negative flow within a channel $\mathcal{C}$ with talk regarding positive flow within its contraposition channel $\mathcal{C}^*$. A further consequence is that there is no need to define the duals of certain concepts, such as pseudo-signal and exception (see below)—e.g. the dual of a pseudo-signal in a channel is simply a pseudo-signal in the contraposition channel. Finally, I will drop the superscript on the classification relation for the rest of the thesis—classification in a channel is always taken to be positive.

I will make use of the following notational simplification throughout the thesis: I will often write $a \mapsto b$ for a connection c such that $signal(c) = a$ and $target(c) = b$; and I will often write $\phi \rightarrow \psi$ for a constraint γ such that $ante(\gamma) = \phi$ and $succ(\gamma) = \psi$. As I stress a number of times, connections and constraints should not be conflated with their endpoints (i.e. the tokens and types they respectively connect). However, this notational simplification will prove to be most convenient and significantly improves clarity of the presentation.

The presence of a classification relation between connections and constraints introduces a characterisation that is not available in links, namely, the characterisation that a given connection is not classified by a particular constraint. This allows *S&B* to model *exceptions* to constraints, as defined below.

2.3.4 Errors: Exceptions and Pseudo-Signals

As I have discussed earlier, a central feature of channel theory is that it incorporates a treatment of "error" or "exception to the general rule". In a previous section, I discussed how pseudo-signals in an information link can lead to error—i.e. if a signal that is not connected to any target is nevertheless used as a basis for inference as if it *had* been connected. The concept of pseudo-signal is also available in a channel.

Definition *Let $\mathcal{C} : A \Rightarrow B$ be a channel. A* pseudo-source *or* pseudo-signal *for a constraint $\gamma \in typ(\mathcal{C})$ is a token $a \in tok(A)$ such that $(a : ante(\gamma))$ holds in A and there is* *no connection* $c \in tok(\mathcal{C})$ *such that* $source(c) = a$.

This concept exactly mirrors the concept of pseudo-signal in an information link. However, a much more interesting possibility for error is provided by the fact that a channel incorporates a classification between connections and constraints. This allows a further distinction to be made, namely, when the source of a connection is classified by the antecedent of some constraint but the connection itself is not classified by that constraint.

Definition *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel. A connection* $c \in tok(\mathcal{C})$ *is an* exception *to a constraint* $\gamma \in typ(\mathcal{C})$ *in* $\mathcal{C}$ *if* $(source(c) : ante(\gamma))$ *holds in* A *and* $(c : \gamma)$ *does not hold in* $\mathcal{C}$.

Given an exception c to a constraint γ, there is a lack of information-flow between the tokens connected by c, even if the source token is classified by the antecedent of the constraint. Since the antecedent condition of the Principle of Harmony does not hold, there is no guarantee that the source and target tokens of the connection are classified by the antecedent and consequent types of the constraint. However, since the source *is* classified by the antecedent, an agent which bases its reasoning on the channel in question (in the manner described in the next section) may come to the conclusion that the target token is also classified by the consequent.

As with an information link, we can define the concept of soundness for a channel: a channel is *sound* just in case it contains no exceptions. Of course, even though a particular connection c is an exception to a given constraint γ, an inference based on c and γ may still happen to lead to a conclusion that actually is supported in the relevant classification. Clearly, this does not justify the inference—any such conclusion is correct simply by chance.

Definition *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel. A connection* $c \in tok(\mathcal{C})$ *is a* weak exception *to a constraint* $\gamma \in tok(\mathcal{C})$ *if* c *is an exception to* γ *and yet* $(target(c) : succ(\gamma))$ *holds in* B. *A* strong exception *is an exception that is not a weak exception.*

Any prediction based on a weak exception is basically correct simply by coincidence. Weak exceptions are not of much interest in the theory of channels, but play a role in the following section, where I investigate the relationship between channels and links.

Exceptions versus Pseudo-Signals

The formal treatment of error in a channel, whether it be via exception or pseudo-signal, localises error—a particular token or connection falling foul of a regularity does not invalidate the regularity, or other information inferred by it.[32] For example,[33] consider a diagrammatic representation of the plan of a house: there is a channel $\mathcal{C}$ between the classification that classifies features of the plan with their position on the diagram and the classification that classifies components of the house with their position in the house. Constraints in this channel have the form something like

$$\textit{feature-of-sort-X-at-posn-Y} \rightarrow \textit{component-of-sort-f(X)-at-posn-g(Y).}$$

Now, suppose the representation of a door on the diagram does not accurately reflect the position of the door in the house—the connection between the representation of the door and the actual door constitutes an exception to the above constraint in $\mathcal{C}$. However, the connections between *other* features of the diagram and their counterparts in the house can safely be used to infer any necessary information regarding the layout of the latter.[34]

A tricky question that arises is the following: when is an error a pseudo-signal and when is it an exception? Basically, this boils down to the question as to when an "abnormal" connection is a connection and when it is something else. The answer is, roughly, "whenever you like (within reason)!". A channel is a classification and, as such, is based on a classifying agent's scheme of individuation. Whether an error corresponds to an exception or to a pseudo-signal depends on how the channel-classification is constructed, and which tokens (i.e. connections) are discriminated by the agent involved.

For example, consider again the *Rey* channel $\mathcal{R}$,[35] involving a thermometer t and patient p. Suppose the conditions involving the use of t with respect to p are somehow abnormal; e.g. assume t has been lying around for some time after being extracted from p. The nurse whose duty it is to monitor the reading of t, and thereby determine p's body temperature, may well be unaware of any problem—he may assume that all is as it should be and makes use of the connection c between t and p as a token of $\mathcal{R}$, leading

[32] This property is a very important one in the semantics of generics, and is discussed further in Chapter 4.

[33] This example is borrowed from *B&S* (1994).

[34] At least, this is theoretically the case. Of course, it could be that the problem with the door is simply a symptom of the whole plan being unreliable.

[35] This example was suggested to me by Jerry Seligman (personal communication).

to a body-temperature estimate that is closer to room temperature than it should be. The alternative view is that the connection c should not be a token of $\mathcal{R}$, but of some other channel $\mathcal{R}'$—constraints in $\mathcal{R}'$ would take a time-delay into account. A nurse that was aware of the delay would therefore not include c in the set of tokens of $\mathcal{R}$—t would therefore be a pseudo-signal in this channel.[36] This shows that one channel's exception can be another channel's pseudo-signal, depending on how the associated connections are individuated.

Throughout this thesis, I will be a little more systematic as to the difference between pseudo-signals and exceptions, particularly when dealing with default reasoning in Chapter 5. Informally speaking, an error will be modelled as a pseudo-signal when there is sufficient information to determine that it actually is an error. For example, if I am told that *opus* is a penguin, then *opus* will be a pseudo-signal in the channel containing the constraint $bird \to flies$. However, if *opus* happens to be a mysterious, magical, flying penguin, then the connection involving *opus* will be an exception to the constraint $penguin \to \neg flies$, in the appropriate channel. This issue is discussed further in Section 5.1.4.[37]

2.3.5 Discussion of Properties

Having presented the basics of channel theory, it is worthwhile spending some time discussing some of the properties of channels and advantages of the theory.

Intensionality of Regularities

A major attraction of channel theory is that it involves minimal commitment as to the constitution of regularities and how they arise. There is an aspect of realism—*B&S* see regularities as somehow being "out there", and that an agent can use these as the basis of inference and other cognitive activity—but there is no attempt to reduce regularities to more fundamental concepts, such as relations between possible worlds. Of course, there is the claim that regularities have structure and that the connections between

[36] Since this nurse may not know enough about the structure of $\mathcal{R}'$—i.e. exactly how the time-delay affects the thermometer-reading—he may well decide that a new reading-event is called for.

[37] Section 4.2.2 describes how the different sorts of errors available in the channel theoretic framework is useful in distinguishing the different sorts of exceptions that need to be catered for in accounts of generic sentences.

tokens plays a fundamental a role as the connections between types. However, this is not the same as providing a reductionist exposition of regularity—the connections at both levels are primitive intensional objects of the theory.

The possible problem with taking a non-reductionist stance with regard to regularities is that it seems to sacrifice explanatory power—i.e. there seems little hope of being able to say anything useful regarding the tasks we may wish to model using channels. However, this is not the case—the concepts that arise from the type-token distinction allow for some very interesting characterisation of various phenomena when channel theory is as a formal modelling tool (e.g. Healey and Vogel (1994) define a channel theoretic model of dialogue that uses the concept of pseudo-signal to characterise certain communication errors). Also, work in subsequent chapters demonstrates how channels not only possess properties that make them particularly attractive as the basis for a semantics of conditionals and generics, but also that powerful logics for conditionals and generics can be defined within the channel theoretic framework. The fact that a powerful system is obtained with little need for any commitment to a particular view of regularity is a very attractive feature.

The non-reductionist stance does have some potential drawbacks, however. When regularities are reduced to relations between other objects, then a logic for applying those generics tends to arise as a result. For example, Morreau (1992a) describes how the problem with the approach that generics are "simply true"—i.e. true in and of themselves, without relation to the behaviour of the associated individuals—is that a logic for using generics to reason defeasibly about the individuals themselves cannot be directly obtained. He uses this argument to support his view of generics as statements about "normal" individuals.[38] Consequently, the logic of generics of Chapter 4 requires the imposition of a *maximal normality* condition in order to obtain a logic for defeasibly reasoning about individuals in Chapter 5. I would argue that the separation between the "truth" of generics and the way in which they are used in default reasoning is a positive thing.

[38] Morreau's work is discussed in some detail in Chapter 4.

Contextuality and Relativism

The channel theoretic model of information flow can be seen as a development of ideas developed in the situation theoretic literature. In particular, the theory of channels can be seen as providing a concrete setting for *conditional constraints*. In situation theory, a conditional constraint involved infons with an *involves* relation, whose appropriate arguments are pairs of situation-types (e.g. (Barwise 1986; Barwise and Perry 1983)). Such an infon is supported by a situation s if s provides a context within which the constraints holds. One view of a channel is that it is just such a situation—i.e. channel theory is a framework in which the structure of constraint-supporting situations is fleshed out.

The set of connections of a channel $\mathcal{C}$ is determined by the context under which the constraints of that channel are appropriate—i.e. it is intended that $tok(\mathcal{C})$ contains only those connections to which the constraints in $typ(\mathcal{C})$ are applicable. For example, if $\mathcal{C}$ is the electric-circuit channel, then connections corresponding to bits of string should not be included in $tok(\mathcal{C})$. Of course, a channel is simply a classification, and it could be that an inappropriate token is included in this classification—such a token is invariably an exception. The sort of information that a channel $\mathcal{C}$ supports is the following, for every token $c \in tok(\mathcal{C})$ and type $\gamma \in typ(\mathcal{C})$:

> *If $source(c)$ is of type $ante(\gamma)$ then (under normal conditions) $target(c)$ is of type $succ(\gamma)$.*

This leads to an interpretation of conditional sentences within the channel theoretic framework, which is the subject of Chapter 3.

The qualification to "normal conditions" in the above conditional is needed to account for the possibility of exception. However, an important feature of channel theory is that the "normal conditions" or "background assumptions" under which a regularity holds are not explicitly represented in a channel, fitting in with recent work in cognitive science which has stressed the impossibility of representing in some complete way all assumptions under which a regularity holds (e.g. see (Suchman 1987; Winograd and Flores 1986)). The background assumptions behind a regularity determine the structure of the channel (e.g. by constraining the set of connections), but are not explicitly represented and therefore cannot be extracted from the channel. This contrasts with, for example, Barwise's (1986) formulation of conditional constraints—Barwise associates a situation-type with each constraint, whereby the situation-type captures the conditions under which the constraint holds. In Chapter 3, I will argue that *some* representation of background conditions of a regularity is necessary if we are to provide an adequate model of conditional reasoning; however, the formal mechanism proposed does not require the assumptions behind a channel $\mathcal{C}$ to be represented within $\mathcal{C}$ itself, nor that $\mathcal{C}$'s assumptions be fully represented anywhere. Rather, these assumptions are represented by the way $\mathcal{C}$ is related to other channels.

The contextual aspect of channels discussed above plays an important role in the theory and its application—in Chapter 4 I argue that it resolves some problematic issues in the semantics of generic sentences. However, there is another important way in which channel theory is relativistic—i.e. the types and tokens, both of classifications and channels, are dependent on some scheme of individuation. While this issue has been raised before, it is important enough to raise once again. This property provides the basis for a sort of *cognitive semantics* (e.g. (Gardenfors 1993)), and has been exploited by Healey and Vogel (1994) in a multi-agent setting to provide an account of misunderstanding in dialogue. Even within a single-agent, multi-perspective[39] setting, this relativism provides a property whereby the types linked by a constraint γ in a channel $\mathcal{C}$ can be either more discriminatory on the world, leading to more robust regularities, or less discriminatory, leading to more efficient regularities. For example, in a normal environment, the frog is happy to use a regularity based on a constraint linking types *flying-black-thing* and *food*; if the world had been different, however, with many non-food flying-black-things in the frog's environment, then the frog may well have evolved in such a way that it was attuned to a regularity linking the types *flying-black-thing-that-buzzes* and *food*.

The trade-off between regularities that are more efficient and regularities that are more robust is a critical one, and issues related to this have plagued workers in AI, especially those working on planning systems, where the relevant problem is known as the *qualification problem* (e.g. (Georgeff 1987b)). The channel theoretic framework developed in the rest of this thesis seems to offer a different approach to this tradeoff, and to the qualification problem in particular. While this issue is not fully investigated in this thesis, the approach is outlined in Section 5.4.

[39] I am using the term "perspective" here in its informal sense, meaning something like "view of the world".

2.3.6 Inference: Links and Channels

I have so far presented the basic elements of *B&S*'s channel theory, but have said very little about the way I intend to use channels as a basis for inference and conditional reasoning, the topic with which this thesis is primarily concerned. An important requirement is that any definition of the concept of inference must allow for the possibility of drawing conclusions which are somehow "wrong". By basing inference on connections, it turns out that the presence of exceptions, as opposed to pseudo-signals, introduces this possibility.

Given a channel $\mathcal{C} : A \Rightarrow B$, it is immediately obvious which connection/constraint pairings should *not* be used for making a prediction regarding classification B, given some information in A; namely, if connection c is an exception to constraint γ, then the pair $\langle c, \gamma \rangle$ should not be used to infer information about the target of c. However, if we are to allow the possibility of error, we need to model the case where it is somehow not known that a particular connection falls out of the domain of classification of an appropriate constraint. Doing so allows an erroneous inference to be made; since each constraint of a channel is applicable to each of its connections, an erroneous prediction is made by taking an exception and inferring that the target token is classified by the succedent of the corresponding constraint.[40]

It should be clear that the above criterion corresponds to (one of) the differences between channels and information links—a link is basically a channel minus the classification between connections and constraints. Of course, a link is also less primitive than a channel—in the latter, the connections and constraints are not identified with their end-points, so there may be more than one connection in a channel that corresponds to a single pair of tokens in the signalling relation of the corresponding link. However, this is only really an issue in the presence of exceptions, and the notion of inference with respect to a link defined below is based on the assumption that the given link is sound. Since this may fail to be the case, it is this assumption that introduces the possibility of error. Before presenting the definition of inference that is used in this thesis,[41] I investigate the relationship between channels and links.

[40] Of course, this only leads to error if the exception is a *strong* exception. If it is a weak exception, the inference is correct, although simply by coincidence.

[41] Actually, the notion of inference introduced below is only really used in Chapter 5, since it is only when reasoning about the properties of individuals (i.e. tokens) that such inference is required.

From Channels to Links

The way I will treat the difference between channels and links in this thesis is roughly as follows: channels support regularities as they actually occur in the world[42] while the information links are the objects on which inference is based. Links and channels are, of course, closely related—inference is based on regularities, and so any (reliable) link will somehow be based on a channel. This is made explicit in the following: given a channel $\mathcal{C}$, we obtain a link based on $\mathcal{C}$ by making connections and constraints extensional and ignoring the classification relation between them.

Definition *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel. The* link $Link(\mathcal{C}) : A \stackrel{\rightarrow}{\rightarrow} B$ defined from $\mathcal{C}$ *has as indicating relation the set of pairs* $\{\langle ante(\gamma), succ(\gamma)\rangle \mid \gamma \in typ(\mathcal{C})\}$ *and as signalling relation the set of pairs* $\{\langle source(c), target(c)\rangle \mid c \in tok(\mathcal{C})\}$.

Clearly, if a channel $\mathcal{C}$ is sound, then $Link(\mathcal{C})$ must also be sound. However, the converse does not necessarily hold—if $\mathcal{C}$ is unsound, but every exception in $\mathcal{C}$ is a weak exception, then $Link(\mathcal{C})$ will be sound.

From Links to Channels

Going in the other direction, from links to channels, is neither as clear-cut nor as important. Since we have taken the stance that channels are somehow more fundamental (i.e. inference is based on regularity; links are based on channels), we do not generally need to go from channels to links. The technical hitch in trying to do so arises because channels have extra structure (namely, a classification relation and more fine-grained intensionality) which cannot be recaptured from links. However, given a link, we can define a *canonical* channel.

Definition *Let* $L : A \stackrel{\rightarrow}{\rightarrow} B$ *be a link. The* canonical channel $Chan(L) : A \Rightarrow B$ defined from L *is defined as follows:*

1. $tok(Chan(L)) = \{\langle a, b\rangle \mid a \stackrel{L}{\mapsto} b\}$;
2. $typ(Chan(L)) = \{\langle \phi, \psi\rangle \mid \phi \rightarrow_L \psi\}$;
3. $(\langle a, b\rangle : \langle \phi, \psi\rangle)$ *holds in* $Chan(L)$ *iff* $(a : \phi)$ *holds in* A *and* $(b : \psi)$ *holds in* B;
4. *the homomorphisms associated with* $Chan(L)$ *are the obvious ones mapping connections and constraints to their endpoints.*

[42] This is subject to the usual relationship with respect to a scheme of individuation.

Seligman (1993) shows that the constructions *Link* and *Chan* completely characterise the relationship between links and channels, under the following provisos: (i) that a channel contains no weak exceptions, and (ii) no distinct connections and constraints are mapped to the same endpoints.[43]

Inference with Links and Channels

The notion of inference I will employ throughout the thesis is the following one: given a link (or set of links) and an initial set of facts, the inferred facts are those which are supported by the signalling and indicating relations of the link in the obvious way. This is made precise as follows.

Definition *Let $L : A \rightrightarrows B$ be a link and Ψ a collection of propositions formed from the tokens and types of A (i.e. Ψ has the form $\{..., (a : \phi), ...\}$, where $s \in tok(A)$ and $\phi \in typ(A)$). The* predicted inferences from Ψ supported by L, *denoted $Consq_L(\Psi)$, are*

$$\{(b : \psi) \mid (a : \phi) \in \Psi, a \overset{L}{\mapsto} b \text{ and } \phi {\rightarrow_L} \psi\}.$$

If $\mathcal{L}$ is a set of links, then I take $Consq_{\mathcal{L}}(\Psi)$ to be $\bigcup\{Consq_L(\Psi) \mid L \in \mathcal{L}\}$.

Although all inference is underpinned by links, I will sometimes abuse this fact and refer to the predicted inferences supported by a *channel* $\mathcal{C}$. What I really mean to refer to in such a case are the inferences supported by $Link(\mathcal{C})$; however, this abuse is sometimes convenient, since it leads to a simplification of the terminology. It should be kept in mind, however, that it is always the links that are used in inference.

The following notation and terminology will prove useful in what follows.

Definition *Let $L : A \rightrightarrows B$ be a link and Ψ a set of propositions.*

The classification induced by L from Ψ *is C such that*

1. $tok(C) = \{b \mid (b : \psi) \in Consq_L(\Psi)$ *for some* $\psi\}$;
2. $typ(C) = \{\psi \mid (b : \psi) \in Consq_L(\Psi)$ *for some* $b\}$;
3. $(b : \psi)$ *holds in* C *iff* $(b : \psi) \in Consq_L(\Psi)$.

Following S&B, C is said to be the projection of Ψ along L. *I write $\Psi \mid\!\!\leadsto_L \Phi$ if $\Phi \subseteq Consq_L(\Psi)$. $\Psi \mid\!\!\leadsto_L \{(b : \psi)\}$ is abbreviated to $\Psi \mid\!\!\leadsto_L (b : \psi)$.*[44]

[43]I.e., it is not the case that $source(c) = source(c')$ and $target(c) = target(c')$ for distinct connections c, c', and similarly for constraints.

[44]I will often drop the subscripts when the link in question can be determined from the context.

Since a link may be unsound, it may be that some predicted consequences are not supported in the succedent classification, even though all the initial propositions are supported in the antecedent classification. For example, suppose we have a channel $\mathcal{C} : A \Rightarrow B$ and proposition $(a : \phi)$, such that $(a : \phi)$ holds in A. Further suppose that $c \in tok(\mathcal{C})$ and $\gamma \in typ(\mathcal{C})$, with $source(c) = a$, $target(c) = b$, $ante(\gamma) = \phi$ and $succ(\gamma) = \psi$, and that c is a strong exception to γ. Then it is the case that $\{(a : \phi)\} \mid\!\!\leadsto (b : \psi)$, but $(b : \psi)$ does not hold in B. However, we do have the following result.

Proposition 2.3.1 *Let $L : A \rightrightarrows B$ be a sound link*[45] *and Ψ a set of propositions such that $(a : \phi)$ holds in A for each $(a : \phi) \in \Psi$. Then the classification C induced by L from Ψ is a subclassification of B.*

Proof Straightforward. $\square$

Inference by way of a channel clearly resembles the use of *modus ponens*. However, I have not defined any notion of a "proof"—i.e. the set of consequences are not closed under application of a set of channels. Instead, it is the set of channels that are closed under some set of operations—this is the topic of the next section.

The concept of inference in channel theory is developed further in later sections of the thesis, particularly in Chapter 5.

2.3.7 Operations on Channels

As defined above, inference proceeds along a single channel or link. However, it is clearly possible for one inferred proposition to lead to another. For this purpose, a number of operations on channels is required. These operations effectively allow a logic based on channel theory to be defined.

The operations defined below are based on those defined by *S&B* (1993), with the parallel composition based on a similar operation defined by Barwise (1993) in the original channel theoretic framework. In that paper, Barwise showed how certain *Principles of Information Flow* could adequately be captured with these operations. In Section A.4.2, I also show that these operations neatly capture the rules of inference of Relevant Logic. I could define operations on links analogous to those on channels; however,

[45] Recall that if $\mathcal{C}$ is a sound channel, then $Link(\mathcal{C})$ is a sound link.

there is no need to do so, since a link arises from each channel, and therefore I take the channels (including the complex ones) as primary.

Connection Graphs

So as to aid in the formalisation of the following operations, the concept of a *connection graph* is needed. This is required for the definition of the necessary operations on connections, ensuring that they posses the necessary properties. The definition follows that of the same concept by *S&B* (1993), with the addition of an extra operation (parallel composition).

Definition *A* connection graph $\mathcal{G}$ *is a graph whose nodes are tokens and whose edges are connections, together with the following operations: identity, contraposition, serial composition and parallel composition. Given a connection* c, $a \overset{c}{\mapsto} b$ *denotes that* c *has source* a *and target* b. *The identity operation provides a connection* $id(a)$ *for each token* a*; the contraposition operation provides a connection* c^* *for each connection* c*; the serial composition operation provides a connection* $(c_1\,;c_2)$ *for each pair of connections* c_1, c_2 *such that* $target(c_1) = source(c_2)$*; the parallel composition operation provides a connection* $(c_1 \parallel c_2)$ *for each pair of connections* c_1, c_2 *such that* $source(c_1) = source(c_2)$ *and* $target(c_1) = target(c_2)$. *These operations satisfy the following conditions:*

1. $a \overset{id(a)}{\mapsto} a$;
2. $a \overset{(c_1\,;c_2)}{\mapsto} b$ *iff there is a token* t *such that* $a \overset{c_1}{\mapsto} t$ *and* $t \overset{c_2}{\mapsto} b$;
3. *if defined,* $(c_1\,;(c_2\,;c_3)) = ((c_1\,;c_2)\,;c_3)$;
4. *if* $a \overset{c}{\mapsto} b$ *then* $(id(a)\,;c) = (c\,;id(b)) = c$;
5. $a \overset{c^*}{\mapsto} b$ *iff* $b \overset{c}{\mapsto} a$;
6. $id(a)^* = id(a)$;
7. *if* $(c_1\,;c_2)$ *is defined then* $(c_1\,;c_2)^* = (c_2^*\,;c_1^*)$;
8. $(c^*)^* = c$;
9. $a \overset{(c_1 \parallel c_2)}{\mapsto} b$ *iff* $a \overset{c_1}{\mapsto} b$ *and* $a \overset{c_2}{\mapsto} b$;
10. *if defined,* $(c_1 \parallel c_2) = (c_2 \parallel c_1)$;
11. *if defined,* $(c_1 \parallel (c_2 \parallel c_3)) = ((c_1 \parallel c_2) \parallel c_3)$;
12. *if* $(c_1 \parallel c_2)$ *is defined,* $(c_1 \parallel c_2)^* = (c_1^* \parallel c_2^*)$.

While I will not mention it explicitly in each of the definitions below, it is assumed that all connections are taken from a connection graph and that the operations on connections satisfy the principles defined above.

Serial Composition

Serial composition was originally introduced by Barwise (1993) as an aid to modelling Dretske's *Xerox Principle,* which is the information-flow equivalent of the logical inference rule of Transitivity.

Definition *Let $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : B \Rightarrow C$ be channels. The* serial composition *of $\mathcal{C}_1$ and $\mathcal{C}_2$ is a channel $(\mathcal{C}_1\,;\mathcal{C}_2) : A \Rightarrow C$ defined as follows:*

1. $tok(\mathcal{C}_1\,;\mathcal{C}_2) = \{(c_1\,;c_2) \mid c_1 \in tok(\mathcal{C}_1), c_2 \in tok(\mathcal{C}_2)$ *and* $target(c_1) = source(c_2)\}$;
2. $typ(\mathcal{C}_1\,;\mathcal{C}_2) = \{\langle\gamma_1,\gamma_2\rangle \mid \gamma_1 \in \mathcal{C}_1, \gamma_2 \in \mathcal{C}_2$ *and* $succ_{\mathcal{C}_1}(\gamma_1) = ante_{\mathcal{C}_2}(\gamma_2)\}$;
3. $(c_1\,;c_2) :^+ \langle\gamma_1,\gamma_2\rangle$ *in* $(\mathcal{C}_1\,;\mathcal{C}_2)$ *iff* $c_1 :^+ \gamma_1$ *in* $\mathcal{C}_1$ *and* $c_2 :^+ \gamma_2$ *in* $\mathcal{C}_2$;

   $(c_1\,;c_2) :^- \langle\gamma_1,\gamma_2\rangle$ *in* $(\mathcal{C}_1\,;\mathcal{C}_2)$ *iff* $c_1 :^- \gamma_1$ *in* $\mathcal{C}_1$ *and* $c_2 :^- \gamma_2$ *in* $\mathcal{C}_2$;
4. $source_{(\mathcal{C}_1\,;\mathcal{C}_2)}(c_1\,;c_2) = source_{\mathcal{C}_1}(c_1)$ *and* $target_{(\mathcal{C}_1\,;\mathcal{C}_2)}(c_1\,;c_2) = target_{\mathcal{C}_2}(c_2)$;
5. $ante_{(\mathcal{C}_1\,;\mathcal{C}_2)}(\langle\gamma_1,\gamma_2\rangle) = ante_{\mathcal{C}_1}(\gamma_1)$ *and* $succ_{(\mathcal{C}_1\,;\mathcal{C}_2)}(\langle\gamma_1,\gamma_2\rangle) = succ_{\mathcal{C}_2}(\gamma_2)$.

Serial composition is most easily illustrated by way of an example (taken from (Barwise and Seligman 1994)). Consider a doorbell, which is attached to a button located on a porch. There are three classifications involved in this example: R involves the bell-token b and the type *ringing* (i.e. R classifies the bell as to whether or not it is ringing); P involves the button-token t and the type *pushed* (i.e. P classifies the button as to whether or not it is being pressed); O involves the porch-token p and the type *occupied* (i.e. O classifies the porch as to whether or not it is occupied). There are two basic channels involved: $\mathcal{C}_1 : R \Rightarrow P$ contains the connection $b \mapsto t$ and constraint *ringing*$\rightarrow$*pushed* (i.e. $\mathcal{C}_1$ involves information-flow from pressed doorbell-buttons to ringing doorbells); $\mathcal{C}_2 : P \Rightarrow O$ contains the connection $t \mapsto p$ and constraint *pushed*$\rightarrow$*occupied* (i.e. $\mathcal{C}_2$ involves information-flow from occupied porches to pressed doorbell-buttons). The serial composition channel $(\mathcal{C}_1\,;\mathcal{C}_2) : R \Rightarrow O$ contains the connection $b \mapsto p$ and constraint *ringing*$\rightarrow$*occupied* (i.e. the composition channel involves information-flow from occupied porches to ringing doorbells). It is the composition channel that allows us, on hearing the ringing doorbell, to deduce that someone is at the door. This example is illustrated in Figure 2.6.

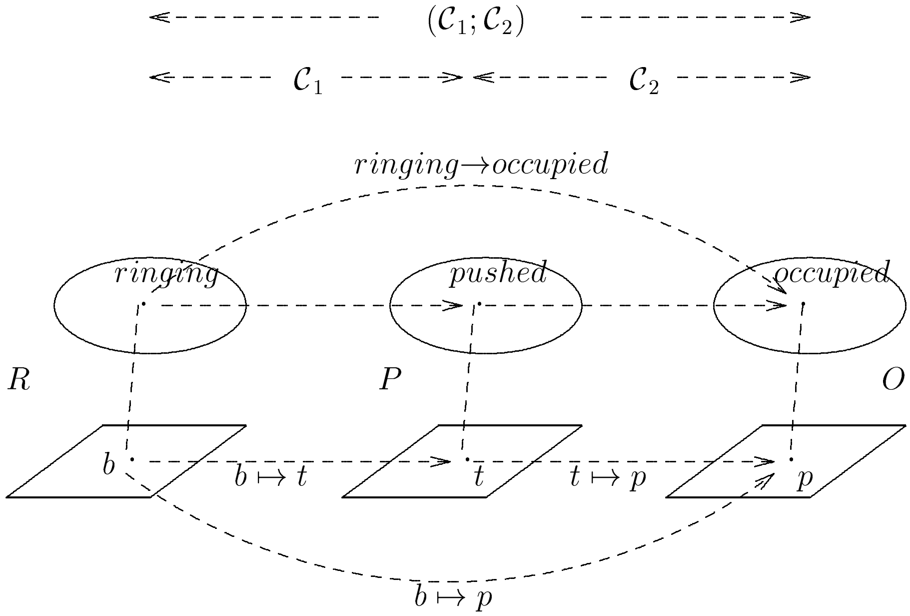

Figure 2.6: Serial Composition: the doorbell example.

S&B (Seligman and Barwise 1993; Seligman 1993) prove various properties of the composition operation, such as associativity and the preservation of exceptions. One of the more interesting results is the following.

Proposition 2.3.2 *(Seligman 1993): $(c_1\,;c_2)$ is an exception to $\langle\gamma_1,\gamma_2\rangle$ in $(\mathcal{C}_1\,;\mathcal{C}_2)$ iff either*

1. *c_1 is an exception to γ_1 in $\mathcal{C}_1$; or*
2. *$(c_1 : \gamma_1)$ holds in $\mathcal{C}_1$ and c_2 is an exception to γ_2 in $\mathcal{C}_2$.*

Parallel Composition

There are several different ways in which a notion of "parallel" composition can be defined. Barwise (1993) required an operation of this form to ensure that his five principles of information flow were satisfied—Barwise's parallel composition specifically addressed the principles of Addition of Information and Exhaustive Cases. For technical reasons, I actually define two separate parallel composition operations, each dealing with the two previously mentioned principles. The *parallel meet* operation deals with constraints that link the conjunction of two types, while the *parallel join* operation deals with constraints that link disjunctions. Recall (from Section 2.2.2) that I assume all classifications to have conjunction and disjunction operation defined over their types.[46] I will sometimes use the term *parallel composition* to refer to both or either of these operations.

Definition *Let $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : A \Rightarrow B$ be channels.*

The parallel meet *of $\mathcal{C}_1$ and $\mathcal{C}_2$ is a channel $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) : A \Rightarrow B$ defined as follows:*

1. $tok(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) = \{(c_1 \parallel c_2) \mid c_1 \in tok(\mathcal{C}_1), c_2 \in tok(\mathcal{C}_2), source(c_1) = source(c_2)$ $\text{and } target(c_1) = target(c_2)\}$;
2. $typ(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) = \{\langle \wedge, \gamma_1, \gamma_2 \rangle \mid \gamma_1 \in typ(\mathcal{C}_1) \text{ and } \gamma_2 \in typ(\mathcal{C}_2)\}$;
3. $(c_1 \parallel c_2) :^+ \langle \wedge, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)$ *iff* $c_1 :^+ \gamma_1$ *in* $\mathcal{C}_1$ *and* $c_2 :^+ \gamma_2$ *in* $\mathcal{C}_2$;

   $(c_1 \parallel c_2) :^- \langle \wedge, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)$ *iff* $c_1 :^- \gamma_1$ *in* $\mathcal{C}_1$ *or* $c_2 :^- \gamma_2$ *in* $\mathcal{C}_2$;
4. $source_{(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)}(c_1 \parallel c_2) = source_{\mathcal{C}_1}(c_1)$ *and* $target_{(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)}(c_1 \parallel c_2) = target_{\mathcal{C}_1}(c_1)$;[47]
5. $ante_{(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)}(\langle \wedge, \gamma_1, \gamma_2 \rangle) = (ante_{\mathcal{C}_1}(\gamma_1) \wedge ante_{\mathcal{C}_2}(\gamma_2))$;

   $succ_{(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)}(\langle \wedge, \gamma_1, \gamma_2 \rangle) = (succ_{\mathcal{C}_1}(\gamma_1) \wedge succ_{\mathcal{C}_2}(\gamma_2))$.[48]

The parallel join *of $\mathcal{C}_1$ and $\mathcal{C}_2$ is a channel $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) : A \Rightarrow B$ defined as follows:*

1. $tok(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) = \{(c_1 \parallel c_2) \mid c_1 \in tok(\mathcal{C}_1), c_2 \in tok(\mathcal{C}_2), source(c_1) = source(c_2)$ $\text{and } target(c_1) = target(c_2)\}$;
2. $typ(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) = \{\langle \vee, \gamma_1, \gamma_2 \rangle \mid \gamma_1 \in typ(\mathcal{C}_1) \text{ and } \gamma_2 \in typ(\mathcal{C}_2)\}$;
3. $(c_1 \parallel c_2) :^+ \langle \vee, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$ *iff* $c_1 :^+ \gamma_1$ *in* $\mathcal{C}_1$ *or* $c_2 :^+ \gamma_2$ *in* $\mathcal{C}_2$;

   $(c_1 \parallel c_2) :^- \langle \vee, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$ *iff* $c_1 :^- \gamma_1$ *in* $\mathcal{C}_1$ *and* $c_2 :^- \gamma_2$ *in* $\mathcal{C}_2$;
4. $source_{(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)}(c_1 \parallel c_2) = source_{\mathcal{C}_1}(c_1)$ *and* $target_{(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)}(c_1 \parallel c_2) = target_{\mathcal{C}_1}(c_1)$;
5. $ante_{(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)}(\langle \vee, \gamma_1, \gamma_2 \rangle) = (ante_{\mathcal{C}_1}(\gamma_1) \vee ante_{\mathcal{C}_2}(\gamma_2))$;

   $succ_{(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)}(\langle \vee, \gamma_1, \gamma_2 \rangle) = (succ_{\mathcal{C}_1}(\gamma_1) \vee succ_{\mathcal{C}_2}(\gamma_2))$.

---

[46] Actually, there is not really a need to assume that the conjunction and disjunction operations are associated with the classifications linked by the channels. An alternative is to have the parallel composition operations link classifications $A \otimes A$ and $B \otimes B$, where $\otimes$ is the sequential conjunction operation on classifications. This possibility is briefly discussed in Appendix A.2.

[47] Recall that $source_{\mathcal{C}_1}(c_1) = source_{\mathcal{C}_2}(c_2)$ and $target_{\mathcal{C}_1}(c_1) = target_{\mathcal{C}_2}(c_2)$

[48] It should be noted that there are two different type-conjunction and type-disjunction operations involved in the right-hand-side of these definitions, depending on whether the operation in question is associated with classification A or classification B. To be precise, I should use different symbols, but I trust that no confusion is caused.

The concept modelled by parallel composition can be illustrated by an example. Consider a television: this involves both an audio channel $\mathcal{A} : T \rightarrow S$ and a visual channel $\mathcal{V} : T \rightarrow S$, where T is a classification of the television and S is a classification of some scene.[49] Given some audio-visual image i, $\mathcal{V}$ carries visual information regarding some scene s while $\mathcal{A}$ carries audio information regarding s. The parallel meet channel $(\mathcal{V} \bigtriangleup \mathcal{A})$ formed from these two channels carries a merger of this information. Note that if either of the channels is not working properly (i.e. $i \mapsto s$ is an exception in one of them), then the parallel join channel can still be informative—i.e. $(i : (\phi \vee \phi'))$ carries the information $(s : (\psi \vee \psi'))$ in $(\mathcal{V} \bigtriangledown \mathcal{A})$ (assuming constraints $\phi \rightarrow \psi$ and $\phi' \rightarrow \psi'$ in the respective channels) so long as at least one of the channels is working normally—whereas the parallel meet channel is informative only when *both* channels are in order. This is reflected in the definition.

It can be shown that the parallel composition of any two channels is itself a channel, as one would hope. This property, as well as properties such as commutativity and associativity of the parallel composition operations, is proved in Appendix A.4.1. As a convenient simplification, I will sometimes write $(\mathcal{C}_1 \parallel \mathcal{C}_2)$ for what I have loosely called "parallel composition". When this is the case, either of the parallel composition operations (i.e. meet or join) can be substituted.

Contraposition

I mentioned above that I do not make use of "negative" information-flow (i.e. from target to source) in the same way as *S&B* do. Instead, I make use of the following *contraposition* operation, which assigns to each channel $\mathcal{C}$ a channel $\mathcal{C}^*$ which models the information-flow in the negative direction.

Definition *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel. The* contraposition *of* $\mathcal{C}$ *is a channel* $\mathcal{C}^* : B \Rightarrow A$ *defined as follows:*

1. $tok(\mathcal{C}^*) = \{c^* \mid c \in \mathcal{C}\}$;
2. $typ(\mathcal{C}^*) = \{\langle -1, \gamma \rangle \mid \gamma \in typ(\mathcal{C})\}$;

[49] It would be preferable to have different classifications involved here—i.e. one classification for the television's visual image and the other for the television's audio image. This can be modelled in the more general setting of parallel composition, as defined in Appendix A.2. Given the current definition of parallel composition, where the two channels involved link the same classifications, we need to express the example in the manner given.

3. $(c^* :^+ \langle -1, \gamma\rangle)$ *holds in* $\mathcal{C}^*$ *iff* $(c :^- \gamma)$ *holds in* $\mathcal{C}$
   $(c^* :^- \langle -1, \gamma\rangle)$ *holds in* $\mathcal{C}^*$ *iff* $(c :^+ \gamma)$ *holds in* $\mathcal{C}$*;*

4. $source_{\mathcal{C}^*}(c^*) = target_{\mathcal{C}}(c)$ *and* $target_{\mathcal{C}^*}(c^*) = source_{\mathcal{C}}(c)$*;*

5. $ante_{\mathcal{C}^*}(\langle -1, \gamma\rangle) = \neg succ_{\mathcal{C}}(\gamma)$ *and* $succ_{\mathcal{C}^*}(\langle -1, \gamma\rangle) = \neg ante_{\mathcal{C}}(\gamma)$.[50]

An exception in a contraposition channel $\mathcal{C}^*$ exactly corresponds to *S&B*'s concept of *exception at the target* in $\mathcal{C}$ (Seligman and Barwise 1993). Further, a channel $\mathcal{C}$ is *sound* in the sense of *S&B* if neither $\mathcal{C}$ nor its contraposition $\mathcal{C}^*$ contains an exception. Throughout the thesis, I will use the terminology I have explicitly introduced—i.e. $\mathcal{C}$ is sound if it does not contain any exceptions; the soundness of $\mathcal{C}^*$ is taken to be a separate issue.

2.3.8 Logical Channels

A particular class of channels that deserves special attention is the class of *logical channels.* Logical or analytical relationships between types constitute one particular form of information-flow: if type ϕ entails type ψ, then a token a being of type ϕ carries the information that a is also of type ψ. I will be especially concerned with logical channels in later chapters, when I present channel theoretic models of conditional and default logics. Since I wish to compare my formal models with previous work, I need to include a "logical" component to the channel theoretic framework to facilitate a more direct comparison. *B&S* (Barwise 1993; Seligman and Barwise 1993) also discuss logical channels—the ideas in this section are based on their observations.

Since a logical channel is concerned with the logical structure of types, any pair of classifications linked by such a channel must be such that the types posses logical structure—i.e. they are *perspectives* in the sense of Section 2.2.2. Since the information-flow is internal to the classification, a logical channel is also assumed to be *reflexive,* in the following sense.

Definition *The signalling relation of a channel* $\mathcal{C}$ *is said to be* reflexive *if, for each connection* $c \in tok(\mathcal{C})$*,* $source(c) = target(c)$*. A channel is said to be* reflexive *if it links the same classification* A *to itself and has a reflexive signalling relation.*

The fact that they are analytic also means that all logical channels are sound—i.e.

[50] As was the case in the definition of parallel composition, there are two different type-negation operations at work here.

contain no exceptions—and are closed under some appropriate set of channel operations. This set of operations is generally assumed to be serial composition, the parallel compositions and contraposition. I refer to this latter condition as the *Logical Closure Constraint:* $\mathcal{L} = (\mathcal{L}\,;\mathcal{L}) = (\mathcal{L} \bigtriangledown \mathcal{L}) = (\mathcal{L} \bigtriangleup \mathcal{L}) = \mathcal{L}^*$.

Axiom *Every* logical channel $\mathcal{L} : A \Rightarrow A$ *is reflexive, sound and satisfies the Logical Closure Constraint.*

Of course, the above definition leaves open many degrees of freedom; this is to be expected, since there are many different logics and a logical channel for each of them.[51] Barwise (1993) presents various examples of logical channels (e.g. classical and intuitionistic logic). For the sake of concreteness, I will tend to assume that logical channels are based on Anderson and Belnap's system **R** of *first-degree relevant entailment* (e.g. (Anderson and Belnap 1975; Dunn 1984)).[52] That is, each logical channel $\mathcal{R} : A \Rightarrow A$ in the remainder of the thesis is (effectively) assumed to support the following: $\gamma \in typ(\mathcal{R})$ iff $ante(\gamma)$ relevantly entails $succ(\gamma)$ in $\mathcal{R}$. (Of course, I should be much more precise here—types aren't necessarily logical formulae. This is spelt out in greater detail in Appendix A.4.2, where I investigate the link between channel operations and rules of inference for **R**.[53]) My reasons for choosing to base logical channels on relevant logic are based on personal preference, but stem from relevant logic's concern with information rather than truth conditions, as mentioned earlier.

2.3.9 Notation and Simplifications

In this chapter, I have presented the concepts from channel theory which I will use in this thesis. *S&B* (Seligman and Barwise 1993; Seligman 1993) have investigated the properties of channels in great detail—for lack of space, I have obviously been unable to present all their results here. For the sake of simplicity, I will make some assumptions throughout the rest of this thesis. These simplifications are by no means necessary and are made simply for the sake of ease of presentation.

[51] Actually, there are many logical channels for each logic since technically there is a separate channel for each classification.

[52] A *first-degree* logic is one with no nested implications.

[53] Restall (1994) has also looked at the link between Channel Theory and Relevant Logic, although his work takes a very different tack to that described in Section A.4.2. Restall views channels as tokens and signalling relations as a ternary relation between tokens. Ternary relations between possible-worlds is the foundation of the Routley-Meyer semantics for relevant logics, and Restall is therefore able to draw interesting links between the two systems. Lemon (1993) has done work along similar lines.

Identification of connections and constraints with their endpoints: I will usually identify connections and constraints in a channel by their endpoints. As discussed above, the tokens and types of channels are primitive objects, and should not be identified with the objects they connect. For example, two tokens may be connected by two different connections, one which gives rise to an exception in a channel and one which does not. However, I will assume this not to be the case throughout the remainder (unless explicitly stated otherwise), allowing me to write $a \mapsto b$ for the connection c such that $source(c) = a$ and $target(c) = b$, and $\phi \rightarrow \psi$ for the constraint γ such that $ante(\gamma) = \phi$ and $succ(\gamma) = \psi$. I will also sometimes refer to the "signalling" and "indicating" relations of a channel—this slight abuse of terminology should cause no problems, given the above assumption.

As an example of the sort of simplification this assumption allows, consider the definition of parallel meet composition, from Section 2.3.7. Identifying connections and constraints with their endpoints allows the parallel meet of channels $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : A \Rightarrow B$ to be defined as follows:

1. $tok(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) = \{a \mapsto b \mid a \mapsto b \in tok(\mathcal{C}_1) \text{ and } a \mapsto b \in tok(\mathcal{C}_2)\}$;
2. $typ(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) = \{(\phi \wedge \phi') \rightarrow (\psi \wedge \psi') \mid \phi \rightarrow \psi \in typ(\mathcal{C}_1) \text{ and } \phi' \rightarrow \psi' \in typ(\mathcal{C}_2)\}$;
3. $(a \mapsto b) :^{+} (\phi \wedge \phi') \rightarrow (\psi \wedge \psi')$ iff $a \mapsto b :^{+} \phi \rightarrow \psi$ in $\mathcal{C}_1$ and $a \mapsto b :^{+} \phi' \rightarrow \psi'$ in $\mathcal{C}_2$;
   $(a \mapsto b) :^{-} (\phi \wedge \phi') \rightarrow (\psi \wedge \psi')$ iff $a \mapsto b :^{-} \phi \rightarrow \psi$ in $\mathcal{C}_1$ or $a \mapsto b :^{-} \phi' \rightarrow \psi'$ in $\mathcal{C}_2$.

While this may not appear a particularly great simplification in this case (mainly because of the cumbersome conjunctive types in the constraints), it does save much in the way of notation in later parts of the text.

Conflation of channels and links: Another abuse I will make is to blur the distinction between channels and links. The identification of connections and constraints with their endpoints is already a step along this path, of course. However, I will sometimes refer to inferences or predictions with respect to a channel $\mathcal{C} : A \Rightarrow B$. Of course, what I really mean in such cases is the set of inferences with respect to the link $Link(\mathcal{C})$. Given the identification of connections and constraints with their endpoints, this effectively involves simply ignoring the classification relation of the channel—i.e. any (strong) exception leads to an incorrect inference with respect to B.

Assumption that signalling relations are reflexive: Some of the work presented in later chapters makes use of the assumption that all signalling relations are *reflexive,* in that $source(c) = target(c)$ for each connection c. This assumption involves great simplification in the presentation—basically, it allows much of the technical work to be performed solely at the level of types—but is also required for a more meaningful comparison of the logics developed in Chapters 3 and 5 with more traditional conditional and default logics. Many of the illustrative examples given in later chapters will reflect this assumption.

Chapter 3

A Channel–Theoretic Model of Conditional Logics

Barwise (1986) has argued that situation theoretic constraints form the semantic content of conditionals—i.e. the semantic content of a statement of the form "*If A then B*" is taken to be a constraint of the form $\phi \rightarrow \psi$, where ϕ and ψ express the content of A and B respectively. Barwise (1993) refines this analysis of conditionals by embedding it in a channel theoretic setting and thereby providing an Austinian analysis—the *descriptive content* of a conditional statement is a type-level regularity, while the *demonstrative content* is a token-level connection.[1] Barwise and Seligman (1994) also suggest that a conditional statement be seen in this light, but with demonstrative content a channel rather than a connection.

In this chapter,[2] I investigate the use of channel theory as a basis for the semantic content of conditionals. My main concern is with the *logic of conditionals* that arises through the use of the channel-operations defined in Section 2.3.7. Various standard logical rules of inference—specifically, *Transitivity*, *Monotonicity* and *Contraposition*—are accepted to be invalid for conditionals (e.g. see (Nute 1980)) and the first task of the evaluation of a logic of conditionals based on the channel-theoretic model is to ensure that these rules are in fact invalidated. I will argue that some extra mechanism is required before the channel operations (as they are defined in Section 2.3.7) are deemed

[1]Informally, the descriptive content of a statement S is the claim made by S, while the demonstrative content of S is the object of which the claim is made.

[2]Some of the work described in this chapter is reported in (Cavedon 1994).

suitable for a logic of conditionals. This mechanism will take the form of a *SubChannel relation*, which more or less orders channels with respect to how "reliable" they are—i.e. a channel that is less reliable, in that it contains more exceptions, is a subchannel of the finer-grained channel. The hierarchy of channels that results can be seen as *encoding* the background assumptions of constraints and regularities, without necessarily *explicitly representing* such assumptions. The channel operations are redefined to take into account a hierarchy of channels against which they are evaluated.

The logic arising from the revised channel operations can be summarised by the following rules, which for convenience are written in a notation roughly corresponding to natural deduction rules (i.e. the antecedent of a rule is written above the line and the consequent is written below it). In the following, the $\mathcal{C} : \phi \rightarrow \psi$ denotes the assertion that the conditional represented by $\phi \rightarrow \psi$ is supported by the channel $\mathcal{C}$.[3]

$$\frac{\mathcal{C}_1 : \phi \rightarrow \psi \quad \mathcal{C}_2 : \psi \rightarrow \tau}{(\mathcal{C}_1 \,;\, \mathcal{C}_2) : \phi \rightarrow \tau}$$ if the background conditions of $\phi \rightarrow \psi$ are compatible with those of $\psi \rightarrow \tau$

$$\frac{\mathcal{C}_1 : \phi_1 \rightarrow \psi_1 \quad \mathcal{C}_2 : \phi_2 \rightarrow \psi_2}{(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) : (\phi_1 \wedge \phi_2) \rightarrow (\psi_1 \wedge \psi_2)}$$ if the background conditions of $\phi_1 \rightarrow \psi_1$ are compatible with those of $\phi_2 \rightarrow \psi_2$

$$\frac{\mathcal{C}_1 : \phi_1 \rightarrow \psi_1 \quad \mathcal{C}_2 : \phi_2 \rightarrow \psi_2}{(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) : (\phi_1 \vee \phi_2) \rightarrow (\psi_1 \vee \psi_2)}$$ without restriction (this reflects the usual acceptance of Weakening as valid)

$$\frac{\mathcal{C}_1 : \phi \rightarrow \psi}{\mathcal{C}_1^* : \neg\psi \rightarrow \neg\phi}$$ if the background conditions of $\phi \rightarrow \psi$ are compatible with $\neg\psi$

The background conditions referred to in the above rules are, as mentioned, captured by the use of a given hierarchy of channels. The resulting logic of conditionals is shown to possess many desirable qualities—in particular, the unwanted rules of Transitivity, Monotonicity and Contraposition are invalidated. On defining various constraints that any given hierarchy is expected to satisfy, it is shown that much of the power of more traditional conditional logics can be retained. Importantly, these constraints can be motivated independently of their use in obtaining a powerful logic of conditionals. The above rules of inference (and the channel-hierarchy mechanism) also form the basis of logics of generics and default reasoning in later chapters of the thesis.

[3] As well as the rules given here, the logic of conditionals also supports the first-degree entailments of the relevant logic **R**.

3.1 A Simple Semantic Model of Conditionals

The channel-theoretic semantic model of conditionals I propose here is very much in line with the ideas of Barwise and Seligman. Basically, a conditional statement involves an assertion that a channel $\mathcal{C}$ supports some sort of "conditional information". This information is related to the internal structure of $\mathcal{C}$, namely, that $\mathcal{C}$ contains some connection c and regularity γ. Barwise's (1986) analysis relates a conditional statement to a constraint—the assertion expressed by a conditional lacks any demonstrative content in this analysis. Barwise (1993) himself later suggests that the demonstrative content of a conditional is a connection between tokens, while Barwise and Seligman (1994) intimate that it is a channel. For reasons discussed below—mainly to do with the analysis of nested conditionals—I will adopt this last stance.

The most urgent issue in a satisfactory logical analysis of conditionals is the problem of counterfactuals: a counterfactual may be true or false, even though its premise is false. For example, the statement

> "If it is raining, I will get wet."

may be true even when it is not raining. As is well known, this property makes the material conditional eminently unsuitable for modelling the conditional connective of such sentences—a more intensional treatment of the connective is required. The standard way of resolving this problem is to resort to possible-worlds and the concept of "nearest-world functions" (e.g. (Stalnaker 1968; Lewis 1973)). Under this treatment, the truth of a logical conditional $A \rightarrow B$ (with respect to a world w) is determined by evaluating B in the "nearest" world w' to w such that A is true in w'. For the example above, the conditional is true in w if the nearest world w' to w in which it is raining is such that I am wet in w'. A specific logic of conditionals arises by imposing various conditions on the nearest-world function f.[4]

Barwise's (1986) analysis of conditionals avoids the potential problem arising from counterfactuals because the truth of a conditional (in his analysis) does not depend on the truth of its constituent statements, but only on the factuality of a constraint. The semantic analysis defined below has the same property, with the extra proviso that, following *B&S*, the content of a conditional is relativised to a channel and also involves

[4]A more detailed presentation of possible-worlds conditional logics is given in Section 3.2.1. See also the survey articles by Nute (1980, 1984).

a connection between tokens. Other properties of conditionals, such as their context-dependency,[5] are also immediately obtained due to the relativistic nature of channel theory.

I will note at this point that there are certain subclasses of conditionals that I consider outside the scope of the analysis of this chapter. Barwise (1986) discusses the partitioning of conditionals into *specific* and *general* conditionals. An example of the latter is the statement

"If it is wet, then the footpaths will be slippery."

This sentence can be seen as making a statement about footpaths in general, rather than about a specific case. While such conditionals could certainly be covered under the semantic analysis below, for convenience I prefer to view them as generic sentences (which they are) and have them fall within the model of generics presented in Chapter 4. Finally, a particular form of conditional is that involving the "even if" construct: e.g.

"Even if it rains, I won't get wet."

Such sentences definitely fall outside the scope of the following analysis, but are discussed in Section 3.1.3.

3.1.1 The Semantics of Conditionals

In situation theory, the content of a declarative utterance is an Austinian proposition, e.g. the claim that a certain situation is of a certain type. This idea can be extended to conditional sentences with channels playing the role of situations—a conditional sentence is taken to be a particular type of declarative sentence, one that makes a claim about a channel $\mathcal{C}$. The claim that such a sentence makes about $\mathcal{C}$ concerns the internal structure of $\mathcal{C}$, namely, that $\mathcal{C}$ supports certain information flow. More precisely, a conditional sentence asserts that a certain channel $\mathcal{C} : A \Rightarrow B$ contains a connection $s \mapsto s'$ amongst its tokens and a constraint $\phi \rightarrow \psi$ amongst its types. Of course, the conditional in question could be a counterfactual. In this case, $(s : \phi)$ does not hold in A. However, $s \mapsto s'$ being a connection of $\mathcal{C}$ and $\phi \rightarrow \psi$ being a constraint in $\mathcal{C}$ in no way depends on whether $(s : \phi)$ holds or not.

[5]That is, conditionals seem to implicitly assume some set of unstated background conditions. This issue assumes prominence below.

In general, the content of a conditional sentence is taken to be an Austinian proposition of the form $(\mathcal{C} : \Pi)$, where $\mathcal{C}$ is a channel and Π is a *conditional fact,* defined below. There are assumed to be classifications that classify channels by conditional facts.

Definition *A* conditional fact *(appropriate for classifying a channel $\mathcal{C}$) is a type involving Austinian propositions formed from the tokens and types of the classifications linked by $\mathcal{C}$: in particular, a conditional fact Π will be written $\langle \Rightarrow, \Phi, \Psi \rangle$, where Φ and Ψ are Austinian propositions. If $\mathcal{K}$ is any classification whose tokens are channels and whose types are conditional facts, then for $\mathcal{C} \in tok(\mathcal{K})$ and $\langle \Rightarrow, (s : \phi), (s' : \psi) \rangle \in typ(\mathcal{K})$,*

$$(\mathcal{C} : \langle \Rightarrow, (s : \phi), (s' : \psi) \rangle) \text{ holds in } \mathcal{K} \text{ iff } s \mapsto s' \in tok(\mathcal{C}) \text{ and } \phi \rightarrow \psi \in typ(\mathcal{C}).$$

The fact that a channel $\mathcal{C}$ is of type $\langle \Rightarrow, \Phi, \Psi \rangle$ can be read as "if Φ holds, then (if conditions are normal)[6] it carries the information Ψ, via $\mathcal{C}$". It should be clear that this conditional information (i.e. that $\mathcal{C}$ is of a certain type) holds regardless of whether Φ (and therefore Ψ) itself holds in the pertinent classification—it simply depends on the internal structure of $\mathcal{C}$. Note that $\mathcal{K}$, being a classification, can be one of the classifications linked by a given channel. This allows a natural interpretation of nested conditionals—the content of a nested conditional is a proposition concerning a channel $\mathcal{C} : A \Rightarrow \mathcal{K}$.

As a simple example, consider the following conditional sentence.

"If the doorbell is ringing, then there is someone on the porch."

This sentence asserts a proposition of the form $(\mathcal{C} : \langle \Rightarrow, (b : \mathit{ringing}), (p : \mathit{occupied}) \rangle)$, where b denotes the doorbell in question, the type *ringing* holds of b just in case b is ringing, p denotes the porch in question, and *occupied* holds of p just in case there is someone on p. This proposition holds whether or not $(b : \mathit{ringing})$ holds. The important point here is that there is a regularity between ringing doorbells and occupied porches (i.e. there is a constraint $\mathit{ringing} \rightarrow \mathit{occupied}$ in $typ(\mathcal{C})$), and the connection between the doorbell and porch in question falls within the domain of this regularity (i.e. there is a connection $b \mapsto p$ in $tok(\mathcal{C})$).

As I mentioned above, a channel effectively plays an analogous role to that of a situation that supports a declarative sentence. An issue that has been raised (via personal communication) as a possible problem is that it is not clear *which* channel should support a given conditional sentence. In the doorbell example, for instance,

[6] The qualification to "normal conditions" is needed since there may be an exception involved.

there may be more than one channel that contains the connection $b \mapsto p$ and constraint *ringing*$\rightarrow$*occupied*—which is the "appropriate" channel? The appropriate channel is, of course, the one that the speaker is describing when she utters the conditional sentence—i.e. there is nothing about channel theory which determines which channel is the most appropriate, just as situation semantics does not determine which is *the* situation described by the sentence "Fido is in front of the fire" (e.g. is the described situation the fire-vicinity situation, the living room situation, the house situation, etc.).[7] What channel theory (and the model of this chapter in particular) is concerned with is: given a channel $\mathcal{C}$ about which certain information is known, what other information can be determined regarding $\mathcal{C}$ (and other channels related to it).

It should be noted that not all conditionals are such obvious instances of general regularities as the one above involving the doorbell. In fact, to account for some conditionals one may sometimes need to take a rather generous view of what constitutes a "regularity". For some conditionals, the pertinent regularity isn't so easily extracted from the conditional sentence. Some potentially problematic examples of conditionals, particular those based on correlation, are discussed in Section 3.1.3 and shown, I believe, to be adequately covered by the semantic notion of conditionals provided by the above definition. By taking the notion of regularity to be the basis of the meaning of conditionals, the assertability of a conditional is based on whether or not there is a connection between the antecedent and consequent of the conditional, thereby ensuring the intensionality of the conditional operator. The semantic analysis of conditionals provided above is clearly a simplistic one; however, it is sufficient for the main purposes of this chapter, which is to provide a channel theoretical logic of conditionals.

3.1.2 Properties of the Channel Theoretic Semantics

The above simple analysis of conditionals directly inherits several properties from the channel theoretic model of information flow. For example, a suitably intensional model of

[7]The fact that a hearer may take an utterance to be about a different situation/token than that intended by the speaker is a starting point for Healey and Vogel's (1994) situation theoretic model of miscommunication in dialogue. A similar model involving "mistaken channel" would be most interesting in the setting of generics, the topic of Chapter 4, since generics are used to make default inferences about individuals. In particular, a speaker and hearer may come to make a different default inference regarding some individual due to them disagreeing as to the nature of the channel supporting the uttered generic sentence. This seems to be a most interesting topic for future research and is outlined in a bit more detail in Section 6.2.

conditionals is obtained, one that adequately models counterfactuals, without a need to introduce possible-worlds or a "nearest-world" selection function. Instead, a conditional is supported by the existence of informational constraints, both at the level of types and tokens, thereby rendering the truth of the conditional itself independent of the truth of its constituent components. However, most importantly to the current enterprise, it is still possible to extract a viable logic for reasoning about conditionals without reducing such connections and regularities to more primitive entities. The contextual aspect provided by a channel as demonstrative content can also be used to explain various phenomena—in particular, the contextual component is used to account for the invalidity of Transitivity, Monotonicity and Contraposition.

The Roles of Regularity and Context

Barwise (1986) examines a number of problematic examples from the literature on conditionals and describes how the information-flow analysis deals with each of these examples. In particular, requiring a conditional to be based on a regularity—rather than attempting to reduce regularities to more primitive concepts—and the role of background conditions seem to account for the various problems raised. For example, consider the following pair of sentences, the first of which is uttered by Tweedledee and the second by Tweedledum:[8]

> "If we had tossed the coin, it would have come up heads."
> "If we had tossed the coin, it would have come up tails."

Under Stalnaker's analysis, each of these conditionals is indeterminate in truth value but Barwise claims them to be false (and I agree with Barwise). The information-flow analysis falsifies both conditionals because of the lack of a regularity underlying either of them. Similarly, conditionals with valid, or necessarily true, consequents are not necessarily supported by the information-flow analysis. For example:

> "If the moon is made of green cheese then $2 + 2 = 4$."

is not supported in the channel-theoretic analysis but is validated by any standard possible-worlds analysis.[9] Finally, the information-flow analysis means that the same

[8] This example is due to Stalnaker (1987).
[9] An exception is Mares and Fuhrmann's (1993) conditional logic based on a Routley-Meyer semantics for relevant logic.

framework is used in the interpretation of both indicative and subjunctive conditionals.[10]

Context and background assumptions also play a crucial role in the interpretation of conditionals since there is usually some assumed condition which is left unstated in a conditional's antecedent. For example, the following conditional is acceptable

> "If I strike the match then it will light."

even though it does not refer to the fact that the match must be dry, that it must be struck in the presence of oxygen, or any other of the multitude of possible assumptions on which the conditional rests. Further, the use of context helps to explain how some conflicting conditionals seem to be simultaneously acceptable. For example, Barwise describes the example of two observers of a card game who respectively assert the following:

> "If Pete calls then he will win."
> "If Pete calls then he will lose."

The first observer bases his claim on the knowledge that Pete is cheating and knows his opponent's cards and will therefore only call if he has a winning hand. The second observer does not knows this but has seen the respective hands (and knows that Pete's hand will lose). In this example, the conflicting conditionals are supported by different channels—the first is supported by a channel containing the regularity that (when he's cheating) Pete only calls when he has a winning hand while the second is supported by a channel containing the regularity relating winning and losing hands with respect to the rules of poker.[11]

Nested Conditionals

For the most part, I will be ignoring nested conditionals below.[12] However, it is most important to show that an inability to interpret nested conditionals is not an inherent shortcoming of the channel theoretic analysis. In fact, since any channel is itself a classification, nested conditionals are handled without any extension to the current

[10] There may be problems with the channel theoretic analysis of subjunctives if one were to assume an extreme realist stance to token-level connections. However, I do not adopt any such stance.

[11] The role of context in the interpretation of conflicting generics is further discussed in Chapter 4.

[12] In particular, the channel theoretic logic of conditionals defined below assumes that conditionals are not nested. Appendix B.1 briefly discusses how the logic could be extended so as to handle nested conditionals.

framework.

Jackson (1987) claims that the only meaningful nested conditionals are those of the form "if X then if Y then Z"—i.e. a conditional whose consequent is itself a conditional. For simplicity, I will restrict attention to nesting of this form. The definition in Section 3.1.1 defined the content of a conditional sentence to be a proposition of the form $(\mathcal{C} : \Pi)$, where $\mathcal{C}$ is a channel and Π a conditional fact. A conditional fact relates two Austinian propositions and one of these could well be the content of some conditional sentence—i.e. the token of the proposition would be a channel and the type a conditional fact. For example, consider the following conditional sentence:

> "If Pete doesn't have an umbrella then if it rains he'll get wet."

This sentence describes a channel $\mathcal{C}$, claiming that it contains a connection $s \mapsto \mathcal{C}'$ and constraint $\neg umbrella \rightarrow \Pi'$, where s is a situation, $umbrella$ classifies those situations in which Pete has an umbrella, $\mathcal{C}'$ is a channel and Π' is a conditional fact. The channel $\mathcal{C}'$ is itself the demonstrative content of the inner conditional—i.e.

> "... if it rains he'll get wet."

This particular conditional claims of $\mathcal{C}'$ that it contains a connection $s \mapsto s$ and constraint $raining \rightarrow wet$, where $raining$ classifies those situations in which it is raining and wet classifies those situations in which Pete is wet. It should hopefully be clear that this interpretation does not require any added machinery beyond that provided by the analysis above.

An interesting issue concerning the choice of demonstrative content for conditional sentences revealed itself when I first investigated the analysis of nested conditionals. Barwise and Seligman (1994) suggest that the demonstrative content of a conditional is a connection between tokens—i.e. the particular connection that supports the (potential) information-flow described by the conditional—and the descriptive content is the corresponding constraint. However, due to the Principle of Harmony, this imposes unrealistic requirements on the antecedent and consequents of the inner conditional. For example, consider again the previous conditional involving Pete and suppose that the demonstrative content of this conditional were the connection $s \mapsto c$, where c is itself the connection $s \mapsto s$. Now suppose that Pete did indeed forget his umbrella (i.e. $(s : \neg umbrella)$ holds) and further suppose that $(s \mapsto c : \neg umbrella \rightarrow \gamma)$ holds (where γ is the constraint $raining \rightarrow wet$). By the Principle of Harmony, it *must* be the case that $(c : \gamma)$ holds which further means that (again by the Principle of Harmony) it *must* be the case that $(s : raining)$ and $(s : wet)$ both hold. However, this conflicts with intuition—if Pete forgot his umbrella then it may still fail to rain in which case Pete wouldn't get wet.[13]

The above problem led to the revision of the semantic content of conditional sentences whereby the demonstrative content is taken to be a channel[14] and the descriptive content is a conditional fact. This seems to provide a more satisfactory interpretation of a nested conditional—in the above example, the fact that Pete forgot his umbrella carries the (conditional) information that if it rains then he'll get wet, without carrying the information that it actually is raining or that he actually does get wet.

3.1.3 Potentially Problematic Conditionals

While my main concern is in defining an adequate logic for conditionals, which in turn is the foundation of the logics of generics and default reasoning of later chapters, I want to briefly discuss some of the potential problems with the channel theoretic view of conditionals presented above. The information-flow analysis of conditionals has several desirable properties but some conditionals do not seem immediately suited to such an analysis. Pollock (1976) distinguishes four classes of conditionals: *simple* conditionals; *necessitation* conditionals; *even-if* conditionals; and *might-be* conditionals. Of these, necessitation conditionals are those requiring some sort of connection between antecedent and consequent, and it is to this class of conditionals that the channel theoretic analysis seems best suited.[15] Might-be conditionals require the use of a modality and I consider them to be outside the scope of the current analysis.

---

[13] One way to avoid the problem is to deny that $(s \mapsto c : \neg umbrella \rightarrow \gamma)$ holds. This would mean that $s \mapsto c$ was an *exception* in the channel supporting the nested conditional, which seems rather unsatisfactory.

[14] This was actually suggested by Jerry Seligman (personal communication) on discussion of the above problem.

[15] Pollock (1976) presents a problematic instance of Transitivity for necessitation conditionals that most conditional logics do not invalidate (e.g. all standard possible-worlds conditional logics validate this pattern of inference). I will discuss this problem, and present a solution to it in the channel-theoretic framework, in Section B.3.

Simple Conditionals

Simple conditionals are those with no link between antecedent and consequent. Some such conditionals seem objectionable, such as those whose consequent is necessarily true. For example, standard analyses make the following conditional true in any model satisfying the laws of arithmetic:

"If the moon is made of green cheese then $2 + 2 = 4$."

Some analyses make a conditional true at any world at which both the antecedent and consequent are (contingently) true—for example, Lewis's (1973) logic of conditionals contains the axiom $(A \wedge B) \supset (A \rightarrow B)$ (where $\supset$ is material implication and $\rightarrow$ is the conditional connective). Recent analyses that use relevant logic as basis for a logic of conditionals (e.g. (Mares and Fuhrmann 1993; Hunter 1980)) invalidate such conditionals, requiring a stronger connection between antecedent and consequent.

Some simple conditionals, however, seem acceptable. For example, suppose I am asked whether Jerry was in Edinburgh at some particular time. I don't know for sure, but I do happen to know that he was in Edinburgh at the time of the last two elections. I may then answer the query with the statement:

"If there was an election on at that time, then Jerry was in Edinburgh."

The coincidence of an election and Jerry's being in Edinburgh is the sort of accidental correlation that *B&S* are quite careful to rule out as the sort of relationship that underlies regularities, which in turn form the basis of channels (as discussed in Section 1.3.2).)[16] Hence, at first blush, it would seem that the above conditional is not supported by *any* channel that is acceptable to *B&S*, even though the conditional itself seems a reasonable one to assert. However, it seems that the use of the past tense in the above statement has an important effect—it seems that we *can* claim there is a regularity between elections in the UK and Jerry being in Edinburgh, but this regularity is limited in scope to those two occasions in the past when this occurred. The channel associated with this regularity contains tokens related to those past events, but not possible future events. This is reflected by the fact that the following conditional does not seem acceptable (given no reason other than coincidence for Jerry to be Edinburgh at the time of an election):

[16] Of course, the coincidence of these events may not be accidental; e.g. Jerry may make the journey to Edinburgh to vote in elections. However, I assume for the sake of argument that the coincidence is just that.

"If there is an election on at that time, then Jerry will be in Edinburgh."

Even-If Conditionals

Even-if conditionals have been explicitly argued by some authors to be in no significant way different to other conditionals (e.g. (Stalnaker 1987)). Other authors have taken the opposite viewpoint, completely excluding even-if conditionals from their analysis (e.g. (Hunter 1980)). Certainly, even-if conditionals seem different to other conditionals—e.g. consider the following:

"Even if it rains, I won't get wet."

I may be able to assert this condition in the knowledge that I have remembered my umbrella. It seems dubious to suggest that there may be a channel, containing a constraint between types *raining* and *not.wet,* supporting such a conditional. However, even-if's seem to involve a different assertion to other conditionals, namely, an assertion that a usual regularity *fails* to hold. For example, given the channel $\mathcal{C}$ containing a constraint $raining \rightarrow wet$, the above even-if conditional seems to assert that some situation s (i.e. the future situation in which I am caught in the rain) is *not* involved in connection in $\mathcal{C}$—i.e. s is a pseudo-signal in $\mathcal{C}$. This distinction seems an interesting one, but I will not pursue it further here.

3.2 Logics of Conditionals

The traditional approach to logics of conditionals, due mainly to Stalnaker (e.g. (Stalnaker 1968)) and Lewis (e.g. (Lewis 1973)), involves a possible-worlds framework, incorporating the concept of "nearest-worlds" selection functions. Barwise (1986) defined a rather different analysis, whereby a conditional statement involves an assertion regarding the factuality of a situation-theoretic constraint. By making such constraints conditional—i.e. by assigning a context, in the form of a type, to each such constraint—Barwise is able to avoid the patterns of reasoning that prove problematic to conditionals. Through the use of the operations on channels defined in Section 2.3.7, one would perhaps hope to be able to similarly define an adequate channel-theoretic logic of conditionals. However, as shown in the following, this is not so without first adding an extra mechanism to the framework.

3.2.1 Stalnaker/Lewis Conditional Logics

By far the most popular approach to the semantics of conditionals involves the use of possible worlds and the concept of a "nearest" world. Stalnaker (1968) was the first to make use of this technique. In Stalnaker's analysis, the truth of a conditional $A{\rightarrow}B$ in a world w depends on the world w' minimally different from w, but making A true. This world is given by a *selection function* f, where $f(A, w)$ is the nearest world to w that validates A. We then have $A{\rightarrow}B$ true in w iff B is true in the world $f(w, A)$.

Not much is said about the selection function f—i.e. what constitutes "nearness"—and this function is allowed to vary across models (i.e. different models may have different selection functions associated with them). However, Stalnaker insists on a number of conditions that any such selection function must satisfy:

1. A is true in $f(A, w)$;
2. if A is true in w, then $f(A, w) = w$;
3. if A is true in $f(B, w)$ and B is true in $f(A, w)$, then $f(A, w) = f(B, w)$.

These conditions ensure that certain axioms and patterns of inference hold in the resulting logic.[17] Transitivity is not an axiom of this logic because the antecedents of two different conditionals may be mapped to different "nearest" worlds by a given selection function. For example, given conditionals $A{\rightarrow}B$ and $B{\rightarrow}C$, we may well have $f(A, w) = w'$ and $f(B, w) = w''$, with B true in w' and C true in w'', thereby making each conditional true in w. However, C is not required to be true in w', i.e. in $f(A, w)$, and so it is not necessarily the case that $A{\rightarrow}C$ is true in w. Similarly, Stalnaker's logic does not validate Strengthening of Antecedents or Contraposition.

Several authors, notably David Lewis, have taken exception to Stalnaker's *uniqueness assumption:* i.e. the assumption that there is a unique "closest" world picked out (for each proposition/world pair) by a selection function. In particular, the following controversial axiom of *conditional excluded middle* follows from this assumption:

CEM: $(A{\rightarrow}B) \vee (A{\rightarrow}\neg B)$

For example, CEM leads to the following conclusion:

"Either: if it is raining then the moon is made of green cheese, or: if it is raining

[17] C.f. the way modal logics are determined by the imposition of certain conditions on the accessibility relation associated with the corresponding Kripkean semantics.

then the moon is not made of green cheese"

Lewis (1973) rejects the uniqueness assumption, leading to his *system of spheres* model, whereby many worlds can be equally "similar" to a given world. I will not present Lewis' semantic model here. However, some of the patterns of inference of his resulting logic are further discussed in Section 3.6.[18] [19]

3.2.2 Situation–Theoretic Analyses

Barwise (1986) proposed an analysis of conditionals based on the situation theoretic model of information-flow. Barwise's analysis can be seen as a precursor to the channel theoretic analysis described above—for Barwise, a conditional statement is a claim regarding an information relationship which, for him, is represented by a situation theoretic constraint.

The main problem that Barwise's simple logic of conditionals needs to address is how Transitivity can, in general, be made to fail. In Barwise's model, Transitivity corresponds to the Xerox property, in that two constraints of the form $\phi \rightarrow \psi$ and $\psi \rightarrow \tau$ can be composed to obtain a constraint of the form $\phi \rightarrow \tau$. However, as in the channel theoretic analysis, a conditional statement corresponds (in general) to a conditional (i.e. fallible) constraint, with which is associated a "background condition". Barwise (1986) represents conditional constraints by making the *involves* relation a three-place relation, relating an antecedent type ϕ, a consequent type ψ and a type B capturing the background conditions under which the constraint holds. Such a constraint is written as follows: $\phi \rightarrow \psi \mid B$, meaning "the constraint $\phi \rightarrow \psi$ holds in situations of type B". The Xerox Principle is restricted to be applicable only to a pair of constraints whose background conditions are identical:

> "If B is fixed, then the resulting two place relation is transitive: if $S_1 \rightarrow S_2 \mid B$ and $S_2 \rightarrow S_3 \mid B$ then $S_1 \rightarrow S_2 \mid B$." (Barwise 1986, p. 122)

[18]Note that Lewis' logic, like most other conditional logics with possible-worlds semantics, validates certain axioms that I find objectionable. In particular, any conditional sentence with a valid consequent or unsatisfiable antecedent is itself valid. This particular problem has been addressed by combining possible-worlds semantics with a Routley-Meyer (1973) semantics for relevant logics, e.g. as done by Mares and Fuhrmann (1993). This issue is further discussed in Section 3.6.

[19]There are other interesting semantic analyses of conditionals that are not based on possible-worlds at all, e.g. Adams' (1975) probabilistic account and Gardenfors' (1988) belief revision model of the Ramsey test.

Barwise's criterion for the cases where Xerox can validly be applied has much in common with the approach taken below, described in Section 3.4. The problem with Barwise's approach, however, is that background assumptions are represented by a single type. This state of affairs, where the all background conditions of a constraint are captured in this way, is an unrealistic one (as has been discussed numerous times above). However, Barwise's information-flow analysis has several attractive properties which are inherited by the channel theoretic analysis.[20]

Wobcke (1995) describes an approach to reasoning about conditionals that uses the tools of the Stalnaker-Lewis conditional logics but which he motivates using situation theoretic concerns. Wobcke's approach is particularly relevant within the current setting because he uses a hierarchy of situations to capture background assumptions behind conditionals in a manner that has some similarities to the approach described below. Like Barwise, Wobcke considers the semantic content of a conditional sentence to be a constraint. However, Wobcke attempts to give an account of constraints using causal links in the form of *abstract plan schemata*, a notion borrowed the AI planning literature. Wobcke defines a situated version of a Stalnaker-Lewis conditional logic, which he calls *SC*. Each sentence in *SC* (conditional or otherwise) is of the form $c : A$, where A is a (standard) propositional sentence and c is interpreted as a situation, modelled as a set of worlds. Letting $[.]$ be the interpretation function, a (situated) conditional sentence $c : A \to B$ is true (in an interpretation) if $f([c], A)$ supports B. This is basically the same interpretation of conditional sentences as given in Stalnaker's logic—in particularly, f picks out a unique closest world—except that selection is with respect to partial situations rather than total worlds. Because of this partiality, *SC* invalidates some of the (from an informational point of view) dubious axioms of Stalnaker-Lewis logics, such as *CEM* and *MOD* (the axiom $A \wedge B \supset A \to B$). However, *SC* does contain weaker (situated) versions of some of these axioms—for example, the following is a theorem of *SC*:

$$c : (A \wedge B) \supset c : (A \to B).$$

An interesting aspect of Wobcke's logic is that he considers the collection of situations associated with any given model as arranged in a hierarchy. He defines how such a hierarchy can be constructed from a given plan schema, which (for the current purposes) can be thought of as an abstract causal rule relating the preconditions of an action-type

[20] Some of these were discussed in Section 3.1.2.

to the intended effect of that action-type. For example,

$$charged(battery) \wedge \neg empty(fueltank) \rightarrow started(car)$$

(partially) describes a *start-car* action.[21] Each plan schema is seen as a situation-type. A hierarchy of schemas is obtained from an initial collection by a construction defined by Wobcke so as to correspond to conditions imposed on the selection function f associated with the model-theoretic semantics of *SC*. For example, given a schema P, a *failure schema* for P is inserted into the hierarchy as a subtype of P. Constructing a failure schema for P involves negating some precondition of P and also negating the effect of P. For example, one failure schema associated with the previous causal rule is the following:

$$\neg charged(battery) \wedge \neg empty(fueltank) \rightarrow \neg started(car).$$

The resulting channel hierarchy of situation types is used as the basis for defining the situation accessibility function f of the model theory—$f(s, A)$ is the most general subsituation of s that satisfies A. If s is the situation supporting the *start-car* schema above (and by Wobcke's definition s thereby supports each of $charged(battery)$, $\neg empty(fueltank)$ and $started(car)$), then $f(s, \neg charged(battery)) = s'$ and s' supports the conditional represented as $\neg charged(battery) \rightarrow \neg started(car)$—i.e. the conditional

"If the battery hadn't been charged than the car wouldn't have started".

Wobcke shows that *SC* invalidates the problematic rules of inference for conditional logics (i.e. Transitivity, Monotonicity and Contraposition) and also claims some interesting properties for it (e.g. subjunctive and indicative conditionals are interpreted within the same framework). However, since (the situation theoretic view of) *SC* involves a reductionist account of information-flow constraints (in terms of the situation-selection function), there are several problematic side-effects—for example, the situated version of Lewis' *MOD* is validated, as is any conditional whose consequent is necessarily true. Also, it is not clear how Wobcke intends to account for conditional sentences that do not discuss actions and their effects. It is of interest that Wobcke's analysis (partially) stemmed from an investigation into the nature of plans (Wobcke, personal communication) and the construction of a general framework for relating conditionals and plans is one of the underlying themes of this thesis.

[21] The relationship between conditionals and planning operators is of interest to the current thesis, as discussed in the Introduction. I outline a channel theoretic model for reasoning about actions and their effects in Section 5.4.3.

3.2.3 Towards a Channel–Theoretic Logic of Conditionals

In Section 2.3.7, I defined a number of operations on channels for constructing new channels from existing ones. These operations were *serial composition, parallel composition*[22] and *contraposition.* These operations form the basis of a "logic" of channels—for example, Barwise (1993) shows how similar channel operations can be used to capture his *Principles of Information Flow* and models various standard logical frameworks.

Given the channel-theoretic semantics of conditionals of Section 3.1.1, my original intention was to define a logic of conditionals based on the channel operations of Section 2.3.7. Consider again the doorbell example from Section 2.3.7 (see Figure 2.6). From the facts that channel $\mathcal{C}_1$ supports the conditional

> "If the doorbell is ringing then the button is pushed."

and $\mathcal{C}_2$ supports

> "If the button is pushed then there is someone at the door."

one can infer that the channel resulting from serial composition supports the conditional

> "If the doorbell is ringing then there is someone at the door."

This result seems perfectly reasonable—although the rule of Transitivity is not generally valid for conditionals, some uses of the rule seem to be unproblematic. However, consider the following pair of conditionals, and the conclusion that is obtained by Transitivity.

> (E1) "If the Lib Dems get the most votes, then the Tories lose the election."
> (E2) "If the Tories lose the election, then Labour will win it."
>
> ---
>
> (E3) "If the Lib Dems get the most votes, then Labour will win the election."

The above example is a classic one for demonstrating the invalidity of the rule of Transitivity for conditionals. Intuitively, the problem with the first two conditionals above is that they seem to be evaluated with respect to different background contexts: the first is true in more or less any context,[23] while the second seems to assume a context in which it is taken that either the Tories or Labour will win.[24] This second

[22] Recall that there are actually two parallel composition operations, but I usually refer to them under the one umbrella term.

[23] Actually, this conditional is not necessarily true in Britain, where elections are not conducted under a system of proportional representation. However, I will gloss over such matters for my current purposes.

[24] This is generally the case in national elections in Britain.

assumption actually conflicts with the premise of the first conditional—i.e. there is a "shift" in background conditions. Barwise (1986) resolves this problem by insisting that the background conditions (modelled as a type) for two constraints must be identical if the factuality of the constraint formed by composition is to follow from the factuality of the two given constraints. This "shift" in background conditions is a problem that the channel theoretic framework must capture if the composition operations are to lead to a satisfactory logic of conditionals.

The context of a regularity is provided by a channel $\mathcal{C}$, in particular, by the connections of $\mathcal{C}$ and the tokens connected by those connections. The expectation is that the channel theoretic view of conditionals can be made to avoid problematic cases of Transitivity via an appeal to this context. For instance, in the election example, we may say that the channel $\mathcal{C}_1$ supporting the conditional (E1) above contains only connections between situations that somehow support the information that it is possible for the Lib Dems to win the election, whereas the channel $\mathcal{C}_2$ supporting (E2) contains only connections between situations that somehow support the information that either Labour or the Tories will win it, i.e. that it is *im*possible for the Lib Dems to win the election in such situations. It may therefore be reasoned that $\mathcal{C}_1$ and $\mathcal{C}_2$ provide mutually conflicting contexts, thereby leading to a channel that is meaningless, or one that fails to support the transitive conditional in some other way.[25]

However, even if the use of Transitivity in the election example can be circumvented in this manner, consider the following example:

(C1) "If there is sugar and diesel-oil in the coffee, then there is sugar in the coffee."
(C2) "If there is sugar in the coffee, then it will taste good."

(C3) "If there is sugar and diesel-oil in the coffee, then it will taste good."

The conditional (C1) expresses an analytical, or logical, relationship—it holds unconditionally and does not depend on any particular context. As such, one expects there to be a channel $\mathcal{C}$ supporting (C1) that is somehow universal—i.e. any connection between appropriate tokens is contained in $\mathcal{C}$. For example, suppose the connections in $\mathcal{C}$ are between cups of coffee; then one would expect there to be a channel $\mathcal{C}$ that supports (C1) such that $\mathcal{C}$ contains the connection $c \mapsto c$ for *any* cup of coffee c. Now, suppose that

[25] Restall (1994) has argued that problematic examples of Transitivity can be explained away in this way.

some particular cup of coffee c makes both (C1) and (C2) true *counterfactuals*: i.e. c contains neither sugar nor diesel-oil, but if c did contain sugar, then it would taste good; and if c did contain both sugar and diesel-oil, then it would contain sugar. Hence, it is certainly possible to envisage channels $\mathcal{C}$ and $\mathcal{C}'$, supporting (C1) and (C2) respectively, such that $c \mapsto c$ is contained in both $tok(\mathcal{C})$ and $tok(\mathcal{C}')$. But then the composition channel $(\mathcal{C}\,;\mathcal{C}')$ will also contain the connection $c \mapsto c$, and this channel supports the conditional (C3), a clearly undesirable result.

The problem here is that when the conditionals are counterfactuals, there is nothing about the connections (of the corresponding channels) that rules out the most straightforward composition of the conditionals. For example, a channel $\mathcal{C}'$ supporting (C2) does not reflect that coffee *cannot* contain diesel-oil and yet taste good—it is simply the case that any cup of coffee covered by this channel *does not* contain diesel-oil. As such, a channel *reflects* the background context of a conditional without explicitly representing it. While this is a desirable feature of channel theory, the upshot is that the channel $\mathcal{C}'$ does not in any way represent the information that a cup of coffee satisfying (C2) *cannot* contain diesel-oil, even though $\mathcal{C}'$ is constrained to be such that any cup of coffee c involved in a connection in $tok(\mathcal{C}')$ is such that c happens to not contain diesel-oil.

It seems that background information pertinent to (C2) must somehow be represented more explicitly than is done so in the channel $\mathcal{C}'$. In Barwise's (1986) model of conditionals, this is done by attaching a type ϕ, which holds only of situations where the cup of coffee in question does not contain diesel-oil (or any other foul-tasting substance), to the constraint representing the (C2). However, this falls back onto the need to explicitly represent all the background conditions of a conditional. The solution proposed in the following section is to evaluate a logic of conditionals with respect to a *hierarchy* of channels, where two channels $\mathcal{C}_1$ and $\mathcal{C}_2$ are related in the hierarchy if $\mathcal{C}_2$ is a "more explicit" version of $\mathcal{C}_1$. For example, $\mathcal{C}_1$ may support the conditional (C2) above, whereas $\mathcal{C}_2$ may support the conditional

(C2$'$) "If there is sugar and no diesel-oil in the coffee, then it will taste good."

The conditional (C2') can be seen as a refinement of (C2), with some of the background conditions (i.e. that the coffee should not contain diesel-oil) moved into the antecedent.[26]

[26]Barwise (1986) states that potential problems arise if one tries to handle background conditions by moving them into the antecedent. In general, this may be so. Moving background conditions into the antecedent of a conditional is just one way to make them explicit—there are other ways, which are discussed below.

Alternatively, the channel $\mathcal{C}_2$ can be seen as a refined, or "more reliable", version of $\mathcal{C}_1$. The important point to notice is that the background conditions of the conditional (C2) are still not in any way explicitly represented in $\mathcal{C}_1$. This is critical if we are not to lose one of the most important properties of channel theory, namely, that there is no attempt to explicitly represent all the background assumptions of a regularity within the channel that provides its contextual component. Instead, the background assumptions are represented in a separate channel (i.e. $\mathcal{C}_2$) and via the relation in which the two channels stand. The next section presents the formal concepts underlying this notion of a hierarchy of channels.

The above discussion has concentrated on the inference rule of Transitivity and the channel operation of serial composition. I mentioned above that other standard rules of inference are accepted to be invalid for conditionals, and the other channel operations also need to be modified. For example, the use of parallel meet composition allows the third conditional below to be obtained from the first two:

> "If the Lib Dems get the most votes, then the Tories lose."
> "If the Tories lose, then Labour wins."
>
> ---
>
> "If the Lib Dems get the most votes and the Tories lose, then the Tories lose and Labour wins."

Similarly, the contraposition operation also requires modification.[27]

3.3 Background Assumptions via the SubChannel Relation

In this section, I present the formal definition of the *subchannel* relation discussed above. Given a collection of channels and a subchannel relation holding over them, we obtain a *channel hierarchy* with respect to which a set of composition operations are defined,

[27] Actually, I know of no convincing counterexample to Contraposition, although it is almost universally accepted to be invalid (Hunter (1980) is an exception). For example, all of Stalnaker's (1987) counter-examples involve *even-if* conditionals, which I have argued require a rather different analysis to standard conditionals. Stalnaker's alternative argument involves showing that, in the presence of Weakening of the Consequent (which is universally accepted as valid), Contraposition reduces to Monotonicity. While it would be an interesting exercise to investigate the consequences of validating Contraposition at the expense of Weakening—in particular, this would result in a neater symmetry since then both parallel compositions operations would require modification—I will not pursue the matter here.

leading to a conditional logic that invalidates the unwanted rules of inference.

3.3.1 The SubChannel Relation

The intention is to define a hierarchy of channels that are increasingly discriminatory. This requires the following definition of a relation between constraints—one constraint is seen as a refinement, or "more discriminatory version", of another if the former has either a stronger antecedent or a weaker consequent. For example, the constraint involved in the content of the conditional "If there is sugar in the coffee then it will taste good" is less discriminatory than that involved in the content of "If there is sugar and no diesel-oil in the coffee then it will taste good".

Definition *Let A' and B' be classifications with orderings $\leq_{A'}$ and $\leq_{B'}$ on their type-sets, and let γ, γ' be types taken from channels $\mathcal{C} : A \Rightarrow B$ and $\mathcal{C}' : A' \Rightarrow B'$, respectively. I write $\gamma \leq_{A',B'} \gamma'$,* wrt $\mathcal{C}$ and $\mathcal{C}'$,[28] *if:*

1. $A \angle A', B \angle B'$,
2. $ante_{\mathcal{C}}(\gamma) \leq_{A'} ante_{\mathcal{C}'}(\gamma')$, *and*
3. $succ_{\mathcal{C}'}(\gamma') \leq_{B'} succ_{\mathcal{C}}(\gamma)$.

Note in this definition that the ordering over the succedent classification is used in the opposite direction to the way the ordering is used over the antecedent classification. For simplicity, I will usually assume that antecedents are strengthened and consequents weakened by the addition of a conjunct or disjunct, respectively. However, this is not required to be the case—strengthening and weakening can be performed with respect to any type-ordering relation $\leq_A$.

I now define what it means for one channel to be a subchannel of another. Intuitively, one channel is a subchannel of another if every constraint γ in the former can be mapped to some constraint γ' in the latter, such that γ is a subconstraint of γ'. The concept of subchannel is defined with respect to a collection $\mathcal{F}$ of functions which map between collections of channel-types. This serves to make the subchannel relation more primitive than it otherwise would be. For example, given two channels $\mathcal{C}$ and $\mathcal{C}'$ which satisfy the first four conditions of the definition, and for each $\gamma \in typ(\mathcal{C})$ some $\gamma' \in typ(\mathcal{C}')$ such

[28] γ is sometimes said to be *subconstraint* of γ'. I will usually drop the subscripts and the reference to the specific channels.

that $\gamma \leq \gamma'$, one does not necessarily want $\mathcal{C}$ to be a subchannel of $\mathcal{C}'$.[29] The collection $\mathcal{F}$ is assumed to contain appropriate identity functions and to be closed under function composition.[30]

Definition *Given a collection $\mathcal{CH}$ of channels and collection $\mathcal{F}$ of functions from channel-types to channel-types, let $\mathcal{C} : A \Rightarrow B$ and $\mathcal{C}' : A' \Rightarrow B'$ be channels in $\mathcal{CH}$ and let $f \in \mathcal{F}$ be a function from $typ(\mathcal{C})$ to $typ(\mathcal{C}')$. $\mathcal{C}$ is a $\leq_{A',B'}$* –subchannel *of $\mathcal{C}'$* wrt f, *written $\mathcal{C} \sqsubseteq_{f,A',B'} \mathcal{C}'$, if*

1. *$A \angle A', B \angle B'$;*
2. *$tok(\mathcal{C}) \subseteq tok(\mathcal{C}')$;*
3. *for all $c \in tok(\mathcal{C})$, $source_{\mathcal{C}}(c) = source_{\mathcal{C}'}(c)$ and $target_{\mathcal{C}}(c) = target_{\mathcal{C}'}(c)$;*
4. *for all $c \in tok(\mathcal{C})$ and $\gamma \in typ(\mathcal{C})$, if $(c : \gamma)$ holds in $\mathcal{C}$ then $(c : f(\gamma))$ holds in $\mathcal{C}'$;*
5. *for all $\gamma \in typ(\mathcal{C})$, $\gamma \leq_{A',B'} f(\gamma)$ wrt $\mathcal{C}$ and $\mathcal{C}'$.*

I write $\mathcal{C} \sqsubseteq_{A',B'} \mathcal{C}'$ if there exists an $f \in \mathcal{F}$ such that $\mathcal{C} \sqsubseteq_{f,A',B'} \mathcal{C}'$.

Aspects of the subchannel relation are displayed diagrammatically in Figure 3.1.

I will sometimes write $\mathcal{C} \sqsubset_{A',B'} \mathcal{C}'$ for $\mathcal{C} \sqsubseteq_{A',B'} \mathcal{C}'$ and $\mathcal{C} \neq \mathcal{C}'$, where equality between channels is taken to be the strongest form possible: i.e. $\mathcal{C} = \mathcal{C}'$ iff $class(\mathcal{C}) = class(\mathcal{C}')$, $left(\mathcal{C}) = left(\mathcal{C}')$, and $right(\mathcal{C}) = right(\mathcal{C}')$. Also, if $\mathcal{C}$ is a subchannel of $\mathcal{C}'$, then $\mathcal{C}'$ is said to be a *superchannel* of $\mathcal{C}$. The subchannel ordering $\sqsubseteq_{A',B'}$ is a preorder (assuming that the type-entailment relations $\leq_{A'}$ and $\leq_{B'}$ are preorders) but, unlike the subclassification relation, not necessarily a partial order. As with all the other definitions, I will usually drop the subscripts A', B' when this causes no confusion.

The concept of ordering propositions in a hierarchy has proved useful in a number of applications in the AI literature. For example, Section 4.1.4 discusses how certain problems with the standard AI default logics (in that they do not satisfy Specificity) can be alleviated by imposing an effective ordering on default rules, requiring that a default rule higher up the hierarchy be satisfied in preference to one lower in the hierarchy when a conflict occurs. Gardenfors (1988) uses the concept of *epistemic entrenchment* to resolve similar conflicts in belief revision—if an epistemic state can be revised in two

[29]For example, the conditions may hold by coincidence—$\mathcal{C}'$ may not actually encode real background assumptions of $\mathcal{C}$. The same effect (in fact, a stronger one) could have been achieved by making the subchannel relation primitive.

[30]This is required to ensure that the subchannel relation is reflexive and transitive.

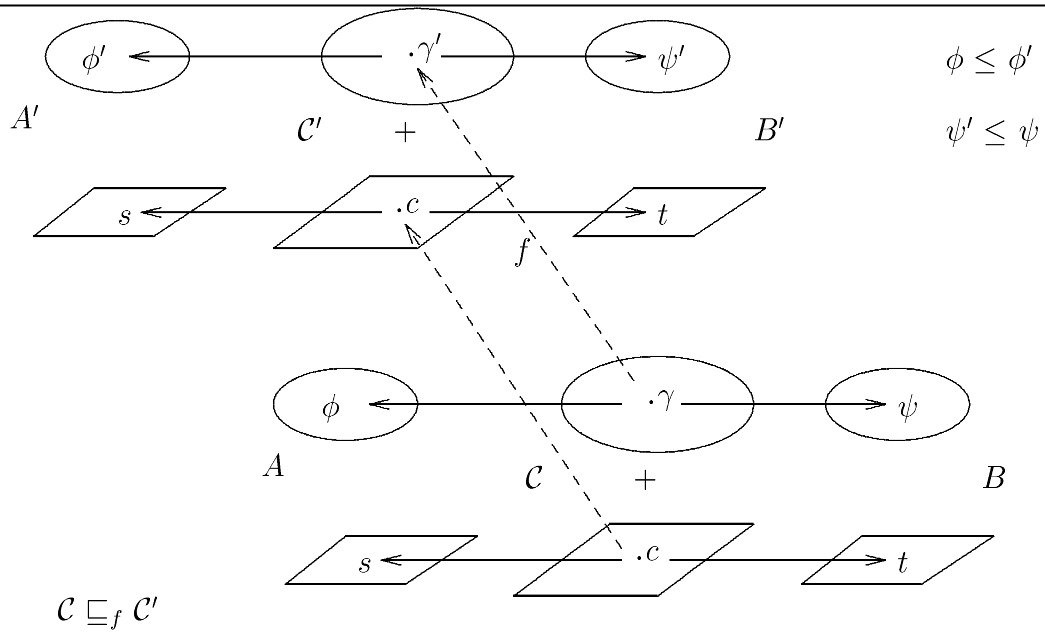

Figure 3.1: Subchannel Relation

different ways so as to achieve the same end (i.e. give up belief of some proposition), then it is revised in such a way that less entrenched formulae are dropped from that state in preference to more entrenched ones. Ryan (1992) defines a general framework for reasoning with ordered theories.

Apart from being purely syntactic (whereas the channel theoretic framework is a semantic one), these other frameworks involving ordered formulae have an important difference to the subchannel ordering in that they allow arbitrary sentences to be related by the ordering, whereas the subchannel ordering can only relate channels whose constraints are related with respect to "type-informedness". This reflects an important difference: the AI uses of orderings are used to address specific problems (in particular, the principle of Specificity in default reasoning) whereas the subchannel ordering is meant to reflect an information-ordering, in that moving up the hierarchy involves moving to a more "informationally robust" channel. This view—that superchannels are more robust versions of their subchannels—is behind the use of the subchannel ordering in representing (implicitly) the background assumptions of fallible regularities (as described below).

3.3.2 Modelling Background Conditions

The above concept of subchannel is intended to capture the notion of implicit background condition in the following way. The background conditions of a constraint $\gamma \in typ(\mathcal{C})$ are made explicit in some constraint $f(\gamma)$, for some f and $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$. This is why I have sometimes talked of $\mathcal{C}'$ being a "refinement" of $\mathcal{C}$—intuitively, $\mathcal{C}'$ contains more discriminatory versions of all the constraints contained in $\mathcal{C}$.[31]

I have discussed earlier that an important feature of channel theory is that the background conditions of a regularity are not explicitly represented. In a given channel hierarchy, this is still the case—the background conditions of a constraint in a channel $\mathcal{C}$ are not represented in $\mathcal{C}$ at all, but may be represented in some channel $\mathcal{C}'$ of which $\mathcal{C}$ is a subchannel. This is very different to representing the conditions within $\mathcal{C}$ itself. Firstly, when a regularity is used in inference, background conditions are left implicit and do not enter into the reasoning process since they are in no way explicated in the channel in question.[32] In particular, a constraint in a channel still counts as a regularity even though there are more refined versions of it in other channels. Secondly, given a constraint γ in a channel $\mathcal{C}$, there need not be *any* channel that makes explicit all of the background conditions of γ. For example, different channels may refine γ in different ways (i.e. different background assumptions associated with the constraints of $\mathcal{C}$ may be represented in different channels, with $\mathcal{C}$ being a subchannel of each of these). Further, two channels $\mathcal{C}_1$ and $\mathcal{C}_2$ with a common subchannel need not have a common superchannel. Even if a channel hierarchy was required to form a lattice, then this could still involve infinitely increasing chains, each channel in a chain involving the addition of some new condition that was left implicit in each of its subchannels.[33]

Hence, a particular channel hierarchy effectively provides another parameter against

[31] In an earlier version of this work, I considered the channel-types in $\mathcal{C}'$ to be the *same* as those in $\mathcal{C}$—i.e. if $\mathcal{C} \sqsubseteq \mathcal{C}'$, then $typ(\mathcal{C}) \subseteq typ(\mathcal{C}')$. However, $\mathcal{C}$ and $\mathcal{C}'$ could have different homomorphisms associated with them, allowing the same type to be mapped to different constraints. For example, γ could be mapped to $(blue.or.green{\rightarrow}sky.coloured)$ by the homomorphisms of $\mathcal{C}$, and to $(blue{\rightarrow}sky.coloured)$ by the homomorphisms of $\mathcal{C}'$. The important difference to the way things are done in the main text is that, on this account, the two channels operate on the *same* type, effectively the same constraint, under two different views or perspectives. This way of viewing things has its attractions, and actually is technically simpler. However, it is dubious whether this fits with *B&S*'s highly relativistic view of classification theory. Hence, I have chosen to take the more general option whereby different channels involve different types/constraints. It would be rather straightforward to modify things to take the other view, with no loss of formal generality.

[32] This issue becomes clearer in Chapter 5, where I outline a methodology for defeasible reasoning in channel theory.

[33] Of course, no such chain could even be envisaged, and there is no need to require it to be.

which certain results (e.g. a logic of conditionals) is relativised, in much the same way as an Austinian proposition is relative to a classification. In the logic of conditionals, the channel hierarchy can be seen as relativising the properties of the logic—i.e. changing the hierarchy, and thereby effectively encoding different background assumptions, modifies the behaviour of the conditional logic.[34] What proves to be important are the properties that *all* hierarchies are constrained to respect, thereby ensuring that the channel theoretic conditional logic displays certain behaviour, no matter the choice of hierarchy (see Section 3.6). Formally, a channel hierarchy could be seen as a structure $\langle \mathcal{CH}, \mathcal{F} \rangle$, consisting of a collection $\mathcal{CH}$ of channels and a collection $\mathcal{F}$ of functions as required by the definition of subchannel. For the most part, I will simply assume that some channel hierarchy is given and that any claims are made relative to it.

I have described the subchannel relation as capturing the notion of increasing discriminability, or "refinement", of channels—given channels $\mathcal{C}$ and $\mathcal{C}'$, where $\mathcal{C} \sqsubseteq \mathcal{C}'$, $\mathcal{C}'$ is somehow a better approximation to its regularities than is $\mathcal{C}$. This is indeed the case if the absence of exceptions is used as a metric of how well a channel captures a regularity.

Proposition 3.3.1 *Let* $\mathcal{C} : A \Rightarrow B$ *and* $\mathcal{C}' : A' \Rightarrow B'$ *be channels, with* $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, *and suppose* $c \in tok(\mathcal{C})$ *and* $\gamma \in typ(\mathcal{C})$*:*

1. *if* c *is an exception to* $f(\gamma)$ *in* $\mathcal{C}'$ *then it must also be an exception to* γ *in* $\mathcal{C}$*;*
2. c *is not necessarily an exception to* $f(\gamma)$ *in* $\mathcal{C}'$ *if it is an exception to* γ *in* $\mathcal{C}$*.*

Proof

1. If c is an exception to $f(\gamma)$, then $c : f(\gamma)$ does not hold in $\mathcal{C}'$—in which case, $c : \gamma$ cannot hold in $\mathcal{C}$ either (by definition of subchannel)—and $source(c) : ante(f(\gamma))$ holds in A'—in which case $source(c) : ante(\gamma)$ must hold in A, since A is a *restriction* of A' (by definition of subchannel). So c is an exception to γ.
2. Even if c is an exception to γ, it may still be the case that $source(c) : ante(f(\gamma))$ does not hold in A', in which case c is not an exception to $f(\gamma)$.

$\square$

Finally, I defined the concept of a logical channel in Section 2.3.8. Logical channels are rather unique in their properties, and one of these is that they are assumed to be

[34]If one so desires, a particular channel hierarchy can be viewed as being associated with a particular agent: the inferences licensed by the resulting conditional logic are then the assertible inferences given the agent's cognitive state.

universally valid—i.e. there are no hidden assumptions or context-dependent factors associated with a logical channel. This is reflected by adding the following condition to the earlier definition of *logical channel*:

Axiom *For any logical channel $\mathcal{L}$, there is no channel $\mathcal{C}$ such that $\mathcal{L} \sqsubseteq_f \mathcal{C}$ or $\mathcal{C} \sqsubseteq_f \mathcal{L}$, except for $\mathcal{L}$ itself.*

That is, logical channels are effectively not involved in the subchannel relation at all.[35] This is not to say that logical channels are not contained in the structure corresponding to a channel hierarchy. In fact, I will generally assume that any hierarchy contains any logical channel required.

3.3.3 Some Assumptions

The rest of this chapter, which deals with a logic of conditionals based on the channel theoretic model, involves the following important simplifications.[36] The first of these is significant, the others simplify the presentation but are easily removed.

The first simplification involves an assumption that the signalling relation of any channel $\mathcal{C}$ is *reflexive* in that, for any connection $c \in tok(\mathcal{C})$, $source(c) = target(c)$. This simplification has a number of consequences. Firstly, it means that most of the complex definitions in what follows involves only types and constraints, rather than tokens and connections as well. In particular, this means that a *type-conflict* relation is powerful enough for the purposes of the logic, rather than requiring a notion of conflicting propositions, allowing simplification to the revised definitions of the channel operations (Section 3.4). A global type-conflict relation is defined below. As an exploitation of this simplification, much of the presentation in the rest of the chapter focuses on constraints, rather than other properties of channels.[37] The removal of this simplification is briefly discussed in Appendix B.2.

The other simplifications mainly involve notation, and are made simply for ease of presentation and readability. This includes the usual conflation between connections and channel-types and the corresponding pairs of tokens and types, respectively. More

[35] An alternative would be to allow logical channels to be related only to other logical channels, with the constraint-mapping function f always being the identity function when restricted to logical channels.

[36] These simplifications are propagated to later chapters.

[37] In particular, all the illustrative figures of this and the following chapters only ever contain the type-level of the relevant channels.

pertinent to the concept of a channel hierarchy is the way in which a superchannel $\mathcal{C}'$ of a given channel $\mathcal{C}$ will be formed usually by adding a type conjunctively to the antecedent of a constraint, or disjunctively to a consequent. As stressed earlier, the earlier definitions are fully general, and this simplification is again made solely for ease of presentation. Finally, the reader will find that some of the definitions below bias towards the assumption that background conditions to a channel are contained only as conjunctions in the antecedents of constraints.[38] However, this is not at all a problem since various conditions imposed on channel hierarchies (such as the *antecedent background constraint* of Section 3.5.1) ensure that such assumptions are completely justified. This point will become clearer below.

A final concept I need is that of *global* type-conflict. This is so I can meaningfully talk about conflicting types when those types are taken from different classifications. Given the simplifying assumption that signalling relations are reflexive and that complex Austinian propositions involve complex types and a single token, type conflict is a strong enough notion to capture the concept of conflicting background assumptions.

Definition *Types $\phi \in typ(A)$ and $\psi \in typ(B)$* globally conflict*, written $\phi \perp \psi$, if there exists a classification C such that $A \sqsubseteq C$, $B \sqsubseteq C$ and $\phi \perp_C \psi$.*

3.4 Channel Operations Relative to a Channel Hierarchy

In Section 3.2.3, I argued that the channel operations as they stand lead to some undesirable patterns of inference being supported by the resulting logic of conditionals, the problem being that the composition operation doesn't adequately account for the implicit background conditions of a regularity. I defined above the concept of a channel hierarchy, within which a constraint's implicit background conditions are represented via the way in which the pertinent channel is related (via the subchannel relation) to others in the hierarchy. I now modify the channel operations so that background conditions, as given by a specific hierarchy, are taken into account, resulting in a system that provides a satisfactory logic of conditionals. Throughout, I assume a particular hierarchy of channels, even though I often leave this parameter unspecified.

[38]For example, the definition of *conditional serial composition* in Section 3.4.1 checks only the antecedents of constraints for conflicting types.

3.4.1 Conditional Serial Composition

Barwise (1986) rules out serial composition of conditional constraints whenever the two constraints in question do not have the same background conditions (represented by a type). The definition of *conditional serial composition* below does more or less the same thing, but with background conditions encoded in a channel hierarchy in the manner described in Section 3.3. The basic idea can be illustrated by the election example. Suppose $\mathcal{C}$ contains the constraint $\gamma = Tories.lose \rightarrow Labour.wins$. The assumption behind this constraint, i.e. that the Lib Dems can't win, is represented in another channel $\mathcal{C}'$, and there is a function f such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, $f(\gamma) = \gamma'$ and $\gamma' = (Tories.lose \wedge Lib.Dems.lose) \rightarrow Labour.wins$. This leads to the following modification of serial composition. Note that I use the same symbol for conditional serial composition as I did for standard serial composition. It may be taken that any instance of serial composition in the rest of the thesis is the one defined below, unless explicitly stated to be otherwise. I trust that this will cause no confusion.

Definition *Let $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : B \Rightarrow C$ be channels, with $\gamma_1 \in typ(\mathcal{C}_1)$ and $\gamma_2 \in typ(\mathcal{C}_2)$. The* conditional serial composition *of $\mathcal{C}_1$ and $\mathcal{C}_2$, written $(\mathcal{C}_1 \,;\, \mathcal{C}_2)$, is identical to the standard serial composition of $\mathcal{C}_1$ and $\mathcal{C}_2$ except that $\langle \gamma_1, \gamma_2 \rangle \in typ(\mathcal{C}_1 \,;\, \mathcal{C}_2)$ only if there do not exist ${\mathcal{C}'}_1, {\mathcal{C}'}_2$ such that*

1. *$\mathcal{C}_1 \sqsubseteq_{f_1} {\mathcal{C}'}_1$, $\mathcal{C}_2 \sqsubseteq_{f_2} {\mathcal{C}'}_2$; and*
2. *$ante_{{\mathcal{C}'}_1}(f_1(\gamma_1)) \perp ante_{{\mathcal{C}'}_2}(f_2(\gamma_2))$.*

The conditions under which conditional serial composition contains a composite constraint is illustrated in Figure 3.2.

By the definition, we need to effectively check all super-channels of the ones in question since no one of them is guaranteed to contain all background assumptions to the relevant constraints (as discussed earlier). Note that whether or not a channel formed by composition contains the composition of two constraints is dependent on a particular channel hierarchy—if the hierarchy does not happen to contain channels with the conflicting antecedents then the composition goes ahead.

As an example of the use of this operation, consider again the doorbell example. Recall that there are channels $\mathcal{C}_1 : B \Rightarrow D$ and $\mathcal{C}_2 : D \Rightarrow P$, with $b \mapsto d \in tok(\mathcal{C}_1)$, $ringing \rightarrow pressed \in typ(\mathcal{C}_1)$, $d \mapsto p \in tok(\mathcal{C}_2)$, and $pressed \rightarrow someone.there \in typ(\mathcal{C}_2)$. The channel $(\mathcal{C}_1 \,;\, \mathcal{C}_2)$ contains the type $ringing \rightarrow someone.there$ unless there are chan-

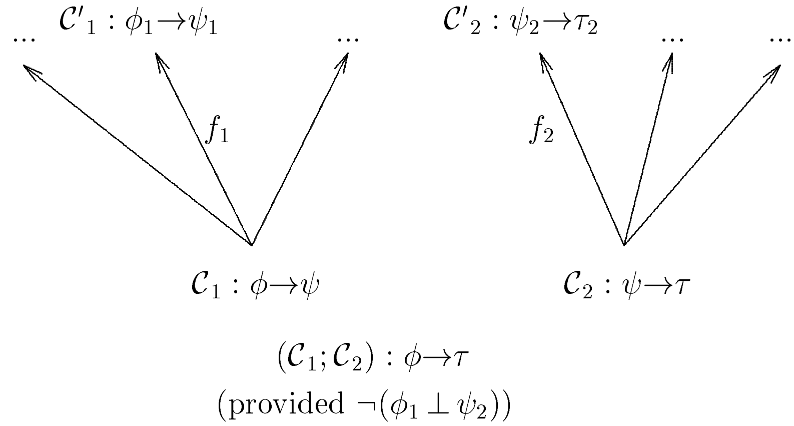

Figure 3.2: Conditions under which the composition channel contains $\phi \rightarrow \psi$.

nels $\mathcal{C}'_1, \mathcal{C}'_2$ containing constraints whose antecedents conflict. Since $\mathcal{C}_1$ and $\mathcal{C}_2$ are associated with different types of information flow, this turns out not to be the case, and the expected composition is allowed.

Now consider the election example. I have already discussed how the channel containing the constraint $Tories.lose \rightarrow Labour.wins$ is related to a channel containing the constraint $(Tories.lose \wedge Lib.Dems.lose) \rightarrow Labour.wins$. Now, one would expect that $Lib.Dems.win \perp Lib.Dems.lose$, which prohibits the composition channel from containing the constraint $Lib.Dems.win \rightarrow Labour.wins$.

So far, I have concentrated on the rule of Transitivity. I now turn briefly to Strengthening of Antecedents. The invalidity of this rule of inference can be demonstrated by the following example:

> "If there is sugar in the coffee, then it tastes good"
>
> ---
>
> "If there is sugar and diesel-oil in the coffee, then it tastes good"

Strengthening of Antecedents follows as a valid pattern of inference from the presence of a logical channel[39] coupled with the rule of Transitivity. For example, the first conditional above can be modelled by a channel $\mathcal{C} : C \Rightarrow C$ containing tokens of the form $c \mapsto c$ and the constraint $sugar \rightarrow good$. Given a logical channel $\mathcal{L}_C : C \Rightarrow C$ containing (amongst

[39] A very simple logical channel will do—for example, one containing constraints of the form $(\phi \wedge \psi) \rightarrow \phi$.

others) the constraint $(sugar \wedge oil) \rightarrow sugar$, then the channel $(\mathcal{L}_C\,;\mathcal{C})$ supports the constraint $(sugar \wedge oil) \rightarrow good$ if the earlier composition operation is used.

The problem with Strengthening of Antecedents is solved by the use of conditional serial composition. Since the background conditions associated with a channel may conflict with the antecedent of the conjunction-eliminating constraint of the logical channel, the undesired inference does not necessarily go through. For example, if channel $\mathcal{C}$ is such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, and $f(sugar \rightarrow good) = (sugar \wedge no.oil) \rightarrow good$,[40] then the conditional serial composition of $\mathcal{L}_C$ and $\mathcal{C}$ does not contain the unwanted constraint. The failure of the conditional composition as demonstrated by this example is illustrated in Figure 3.3.

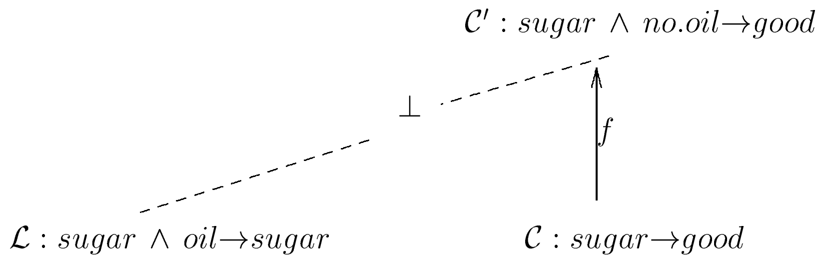

Figure 3.3: Failure of composition for the coffee example.

3.4.2 Conditional Parallel Composition

As discussed earlier, the unqualified use of parallel meet composition can also lead to unsatisfactory inferences. This operation is modified in much the same manner as was serial composition.

Definition *Let $\mathcal{C}_1 : A \Rightarrow B$ and $\mathcal{C}_2 : A \Rightarrow B$ be channels, with $\gamma_1 \in typ(\mathcal{C}_1)$ and $\gamma_2 \in typ(\mathcal{C}_2)$. The* conditional parallel meet composition *of $\mathcal{C}_1$ and $\mathcal{C}_2$, written $\mathcal{C}_1 \triangle \mathcal{C}_2$, is the same as the usual parallel meet composition of $\mathcal{C}_1$ and $\mathcal{C}_2$ except that $\langle \wedge, \gamma_1, \gamma_2 \rangle \in typ(\mathcal{C}_1 \triangle \mathcal{C}_2)$ only if there do not exist $\mathcal{C}'_1, \mathcal{C}'_2$ such that*

1. *$\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C}'_1$, $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C}'_2$; and*
2. *$ante_{\mathcal{C}'_1}(f_1(\gamma_1)) \perp ante_{\mathcal{C}'_2}(f_2(\gamma_2))$.*

[40]I.e. the fact that a given cup of coffee does not contain diesel-oil is a background assumption to the constraint that if it contains sugar it tastes good.

This operation behaves in much the same way as conditional serial composition. The definition of parallel join remains unchanged, reflecting the fact that Weakening of the Consequent is universally accepted as valid in conditional logics. As with conditional serial composition, whenever I refer to "parallel meet", I mean "conditional parallel meet". If I ever refer to "parallel composition", then I mean the the conditional parallel meet operation coupled with the (standard) parallel join operation.

3.4.3 Conditional Contraposition

The modification of the contraposition operation is somewhat more complex than for serial and parallel composition. An initial problem is that, given a constraint that contains a conjunction in the antecedent, after contraposition this becomes a disjunction in the consequent. This seems to conflict with the earlier assumption that any constraint's background conditions are represented in the antecedent of a constraint higher up the channel hierarchy. For example, suppose $\mathcal{C}$ supports some conditional and that $\mathcal{C}'$ supports a conditional with a stronger antecedent—i.e. $\mathcal{C} \sqsubseteq \mathcal{C}'$ and $\mathcal{C}'$ contains (some) background assumption of the conditional supported by $\mathcal{C}$. On contraposing $\mathcal{C}$ to obtain $\mathcal{C}^*$, one would expect the *contraposition* $\mathcal{C}'^*$ of $\mathcal{C}'$ to represent the background conditions of $\mathcal{C}^*$. However, $\mathcal{C}'^*$ will contain a weakened succedent, rather than a strengthened antecedent, which conflicts with the assumption that background assumptions are represented in the antecedent of superchannels.

The solution to the problem is to ensure that even when $\mathcal{C}'^*$ represents background assumptions of $\mathcal{C}^*$, these assumptions are *also* represented in the antecedent of a constraint in some other channel. This resolves the problem entirely. It does not matter that more than one channel represents the implicit assumptions of $\mathcal{C}$, one channel representing the assumptions in the antecedent and the other in the consequent of its constraints. The modified channel operations simply require *some* superchannel to provide a conflicting antecedent for the unwanted composition to be blocked. The formal condition required of the channel hierarchy so as to provide this solution is defined in Section 3.5—for now, I simply continue to assume that any background assumption is represented in the antecedent of a constraint higher up the hierarchy.

The following is the new definition of the contraposition operation, modified to account for assumptions represented in a channel hierarchy. This definition is more complex than the previous ones, and some concepts need to be spelt out somewhat—

this is done below the definition.

Definition *Let $\mathcal{C} : A \Rightarrow B$ be a channel, with $\phi \rightarrow \psi \in typ(\mathcal{C})$. The* conditional contraposition *of $\mathcal{C}$, written $\mathcal{C}^*$, is the same as the usual contraposition of $\mathcal{C}$, except that $\neg\psi \rightarrow \neg\phi \in \mathcal{C}^*$ only if there does not exist $\mathcal{C}' : A' \Rightarrow B'$, where $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, such that*

$$\neg\psi \perp \tau, \textit{ where } \tau \textit{ is the least type in } A' \textit{ such that } ante(f(\phi \rightarrow \psi)) \leq_A (\phi \wedge \tau).$$

The first point I need to address is the concept of "least type". By "least", I simply mean least with respect to the type-ordering $\leq_{A'}$. The constraint $f(\phi \rightarrow \psi)$ is that which $\phi \rightarrow \psi$ is mapped to, so the antecedent of this constraint is the one containing a background assumption of $\phi \rightarrow \psi$. Since τ is the least type that, when added to ϕ, results in a type at least as strong as $ante(f(\phi \rightarrow \psi))$, then τ is effectively the background assumption.[41] (Note that, if I assumed that the only way to strengthen an antecedent was by adding a conjunct, then τ would always be this conjunct.) Hence, the contraposition of the constraint $\phi \rightarrow \psi$ is not contained in $\mathcal{C}^*$ when the antecedent of the contraposed constraint (i.e. the negation of the consequent of the initial constraint) conflicts with any background assumptions of the initial constraint.

Perhaps the best way to illustrate the Conditional Contraposition operation is by way of an example. As discussed above, clear counterexamples to Contraposition are hard to find. However, consider Stalnaker's (1987) argument (demonstrating how Strengthening follows from Contraposition) applied to the usual coffee example. From

> "If the coffee has sugar in it, it tastes good"

the following is obtained by Contraposition:

> "If the coffee tastes bad, there is no sugar in it".

This conditional is supported by some channel $\mathcal{C}_1$. Using parallel composition on a logical channel containing the constraint $\neg oil \rightarrow \neg oil$, the channel $\mathcal{C}$ of Figure 3.4 is obtained. This channel supports the conditional

> "If the coffee tastes bad or there is no oil in it, then either there is no sugar in it or there is no oil in it".

If standard contraposition is applied to $\mathcal{C}$, a channel supporting the following conditional is obtained:

[41] I should really say τ is *a* background assumption, since other background assumptions will be represented in various other channels in the hierarchy.

"If there is sugar and oil in the coffee, then it tastes good and there is oil in it"

which is clearly an unreasonable inference to draw from the conditional I first started with. However, the channel obtained by *conditional contraposition* does not support this conditional. This is due to the presence of the channel $\mathcal{C}'$, of which $\mathcal{C}$ is a subchannel; $\mathcal{C}'$ contains the constraint $\neg good \wedge \neg oil \rightarrow \neg sugar \vee \neg oil$—this constraint was obtained, via parallel composition,[42] from the standard background constraint $\neg good \wedge \neg oil \rightarrow \neg sugar$. Assuming that $(\neg(\neg sugar \vee \neg oil)) \perp (\neg oil)$, the extra condition in the definition of conditional contraposition blocks the unwanted inference.

$$\begin{tikzcd} \mathcal{C}' : (\neg good \wedge \neg oil) \rightarrow (\neg sugar \vee \neg oil) \\ \mathcal{C} : (\neg good \vee \neg oil) \rightarrow (\neg sugar \vee \neg oil) \arrow[u, "f"'] \end{tikzcd}$$

Figure 3.4: Illustration of the Conditional Contraposition operation.

As a final important aside, the following result demonstrates that Contraposition preserves the hierarchy. The proof (of this and all other results of this chapter) is contained in Appendix B.4.

Proposition 3.4.1 *Let $\mathcal{C} : A \Rightarrow B$ and $\mathcal{C}' : A' \Rightarrow B'$ be channels such that $\mathcal{C} \sqsubseteq_{f'} \mathcal{C}'$. Further suppose that $\phi \rightarrow \psi \in typ(\mathcal{C})$ and that $f(\phi \rightarrow \psi) = \phi' \rightarrow \psi'$.*

1. *If $\neg\psi \rightarrow \neg\phi \in typ(\mathcal{C}^*)$ then $\neg\psi' \rightarrow \neg\phi' \in typ(\mathcal{C}'^*)$;*
2. *$\mathcal{C}^* \sqsubseteq \mathcal{C}'^*$.*

3.5 Constraining Channel Hierarchies

In the previous section, I defined new versions of the channel operations, having modified them so as to take account of implicit background assumptions of regularities. The logic that results from these modified operations does not support all inferences analogous

[42] I.e. using parallel join with an identity channel containing the constraint $\neg oil \rightarrow \neg oil$.

to Transitivity, Strengthening of Antecedents, and Contraposition. However, there are no guarantees that this logic satisfies any of the well-accepted patterns of inference of standard conditional logics—since I have specified no conditions on a channel hierarchy other than its existence, *any* implicit condition could be coded into it. For example, a channel supporting the conditional

> "If there is sugar in the coffee and it tastes bad, then it tastes good."

could be inserted in a hierarchy. Such a channel would seem to be meaningless—the hierarchy corresponding to this case seems to be internally inconsistent. However, there is nothing in the current framework that prevents this.

In this section, I define a number of conditions that all channel hierarchies are assumed to satisfy. One such condition removes the possibility of internal inconsistencies of the type described above. Other conditions serve to strengthen the logic by ensuring that certain patterns of inference are supported by any channel hierarchy. This is analogous to the way standard conditional logics are determined by constraints imposed on the corresponding nearest-world selection function (see Section 3.2.1). An interesting property of the following constraints is that they are conditions one may expect *any* channel to satisfy, given the view of channels presented in Chapter 2 and the notion of the channel hierarchy, as is discussed below.

3.5.1 Antecedent Background Constraint

So far, I have assumed that the background assumptions of a regularity supported by a given channel $\mathcal{C}$ are represented in the *antecedent* of constraints of superchannels of $\mathcal{C}$. In some cases, it is more appropriate to weaken a conditional by adding a disjunct to its consequent rather than by adding a conjunct to its antecedent.[43] The following condition ensures that any background assumption associated with a constraint in $\mathcal{C}$ is represented in the antecedent of a constraint in *some* super-channel of $\mathcal{C}$.

Definition Antecedent Background Constraint: *Let $\mathcal{C}$ be a channel with $\phi{\rightarrow}\psi \in typ(\mathcal{C})$, and suppose that $\mathcal{C} \sqsubseteq_{f'} \mathcal{C}'$, where $f'(\phi{\rightarrow}\psi) = \phi'{\rightarrow}\psi'$ and $\psi' \neq \psi$. Let τ be the least type such that $(\psi \vee \tau) \leq \psi'$. Then there exists $\mathcal{C}''$ such that $\mathcal{C} \sqsubseteq_{f''} \mathcal{C}''$ and $f''(\phi{\rightarrow}\psi) = (\phi' \wedge \neg\tau){\rightarrow}\psi$.*[44]

[43] This seems especially true when channel contraposition is available.

[44] It is easily shown that this definition is well-defined, i.e. that $\phi{\rightarrow}\psi \leq f''(\phi{\rightarrow}\psi)$.

This definition requires some explanation. As in the definition of conditional contraposition, the type τ represents a background condition to the constraint $\phi \rightarrow \psi$. The Antecedent Background Constraint simply asserts the existence of a channel $\mathcal{C}''$ that contains a constraint that represents the background assumption τ (as well as any other background assumptions contained in ϕ') in its antecedent, with no background conditions contained in its consequent (this is ensured by the condition that the consequent of $f''(\phi \rightarrow \psi)$ is ψ itself).

To illustrate with an example, consider the standard coffee conditional, supported by some channel $\mathcal{C}$. Let $\mathcal{C}'$ contain a constraint that contains a background assumption in its consequent. This is illustrated by the left half of Figure 3.5. In this example, the type τ is *oil*. The Antecedent Background Constraint requires the channel $\mathcal{C}''$ to be contained in the channel hierarchy, represented by the right side of the figure. Note that the type *oil* has been negated and moved into the antecedent.

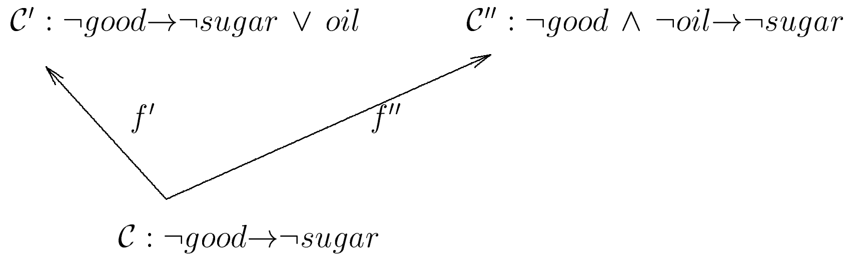

Figure 3.5: Illustration of the Antecedent Background Constraint.

The Antecedent Background Constraint *ensures* that any background conditions of a channel $\mathcal{C}$ are contained in the antecedent of a constraint in some superchannel of $\mathcal{C}$, which allows the simplification of the definition of conditional composition (which only checks antecedents for conflicting background conditions). The alternative is to remove the Antecedent Background Constraint and generalise the definition of the composition operations.

3.5.2 Consequent Background Constraint

Just as the Antecedent Background Constraint ensures that any implicit background assumption is represented in an antecedent of a constraint in some superchannel, the

Consequent Background Constraint below ensures the symmetric condition—that any such assumption is contained in a consequent somewhere in the hierarchy. While the former constraint is needed to ensure that my earlier assumption (i.e. that only antecedents in superchannels need to be checked in the new definitions of the channel operations) is a valid one, it turns out that both Background Constraints are useful in other ways, as shown when I introduce the Parallel and Serial Subchannel Constraints below.

Definition Consequent Background Constraint: *Let $\mathcal{C}$ be a channel with $\phi \rightarrow \psi \in typ(\mathcal{C})$, and suppose that $\mathcal{C} \sqsubseteq_{f'} \mathcal{C}'$, where $f'(\phi \rightarrow \psi) = \phi' \rightarrow \psi'$ and $\phi' \neq \phi$. Let τ be the least type such that $\phi' \leq (\phi \wedge \tau)$. Then there exists $\mathcal{C}''$ such that $\mathcal{C} \sqsubseteq_{f''} \mathcal{C}''$ and $f''(\phi \rightarrow \psi) = \phi \rightarrow (\psi' \vee \neg \tau)$.*[45]

Since the above definition is analogous to the one in Section 3.5.1, I will not explain it any further. Note that the presence of the two constraints together ensure a certain symmetry in a channel hierarchy—any implicit assumption that is contained in a consequent in a superchannel is also contained in an antecedent in some superchannel, and vice versa.

3.5.3 Parallel Subchannel Constraint

The next two constraints concern the interaction between a given channel hierarchy and the channel composition operations.[46] These constraints prove to be important later, when I show that the channel theoretic conditional logic verifies certain patterns of inference. The first constraint concerns the parallel composition operations.[47]

Definition *Let $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}'_1, \mathcal{C}'_2$ be channels, with $\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C}'_1$ and $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C}'_2$. Further suppose $\gamma_1 \in typ(\mathcal{C}_1), \gamma_2 \in typ(\mathcal{C}_2)$ and $\langle \wedge, \gamma_1, \gamma_2 \rangle \in typ(\mathcal{C}_1 \parallel \mathcal{C}_2)$. Then we define $(f_1 \parallel f_2) : typ(\mathcal{C}_1 \parallel \mathcal{C}_2) \rightarrow typ(\mathcal{C}'_1 \parallel \mathcal{C}'_2)$ to be the function that maps $\langle \wedge, \gamma_1, \gamma_2 \rangle$ to $\langle \wedge, f_1(\gamma_1), f_2(\gamma_2) \rangle$, and maps $\langle \vee, \gamma_1, \gamma_2 \rangle$ to $\langle \vee, f_1(\gamma_1), f_2(\gamma_2) \rangle$.*[48]

---

[45]It is easily shown that this definition is well-defined, i.e. that $\phi \rightarrow \psi \leq f''(\phi \rightarrow \psi)$.

[46]I have already shown, in Section 3.4.3, that conditional contraposition interacts smoothly with the channel hierarchy.

[47]Following a convention I mentioned earlier, I will use the symbol '$\parallel$' to represent "parallel composition"—either of the specific parallel composition operations (i.e. meet or join) can be substituted in such cases.

[48]It should be noted that, for technical correctness, the definition requires a proof that the type $\langle \wedge, f_1(\gamma_1), f_2(\gamma_2) \rangle$ is actually contained in channel $(\mathcal{C}'_1 \parallel \mathcal{C}'_2)$—i.e. otherwise it could be that the conditional parallel meet does not contain this constraint because of some conflicting background conditions.

Although this definition may seem rather obscure, its function is fairly straightforward—the defined function maps a type formed by parallel composition of γ_1, γ_2 to the parallel composition of $f_1(\gamma_1), f_2(\gamma_2)$. The definition (confined to the conjunctive parallel types) is illustrated in Figure 3.6.

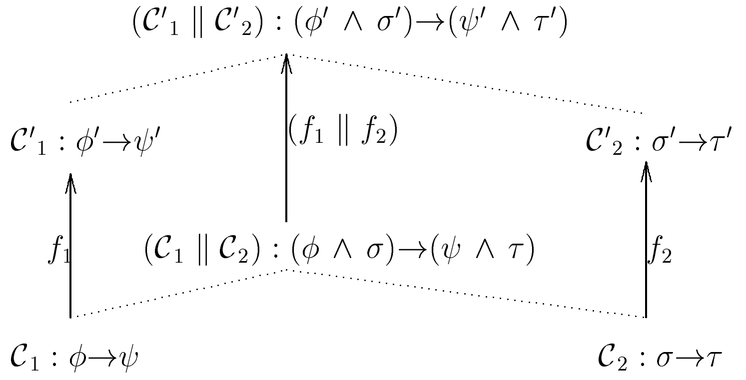

Figure 3.6: Definition of mapping $(f_1 \parallel f_2)$.

The following result shows that the function defined above fits nicely into a channel hierarchy that is closed under parallel composition.

Proposition 3.5.1 *Suppose all the conditions of the above definition hold. Then* $(\mathcal{C}_1 \parallel \mathcal{C}_2) \sqsubseteq_{(f_1 \parallel f_2)} (\mathcal{C}'_1 \parallel \mathcal{C}'_2)$.

This result is an important one. It basically shows that conditional parallel composition "merges" the background conditions of the channels involved—any such assumptions are contained in channels higher up the subchannel ordering, and the channel obtained by parallel composition of the higher channels is itself higher up the ordering than the composition of the lower channels. Of course, not all channels that are higher in the ordering than the channel $(\mathcal{C}_1 \parallel \mathcal{C}_2)$ need themselves be formed from the composition of other channels. However, since those that are formed by composition contain all the background assumptions needed, I will impose the converse of the proposition as a constraint on the channel hierarchy. This proves useful later on.

Definition Parallel Subchannel Constraint: *If* $(\mathcal{C}_1 \parallel \mathcal{C}_2) \sqsubseteq_f \mathcal{C}$, *then there exist* $\mathcal{C}'_1, \mathcal{C}'_2$, *where* $\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C}'_1$ *and* $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C}'_2$, *such that* $\mathcal{C} = (\mathcal{C}'_1 \parallel \mathcal{C}'_2)$ *and* $f = (f_1 \parallel f_2)$.

This is proved in Appendix B.4.

The Parallel Subchannel Constraint ensures that any channel containing background conditions relative to a channel $\mathcal{C}$ formed by parallel composition is itself formed by parallel composition from channels that contain background conditions pertinent to the component channels of $\mathcal{C}$.

The Parallel Subchannel Constraint is illustrated in Figure 3.7.

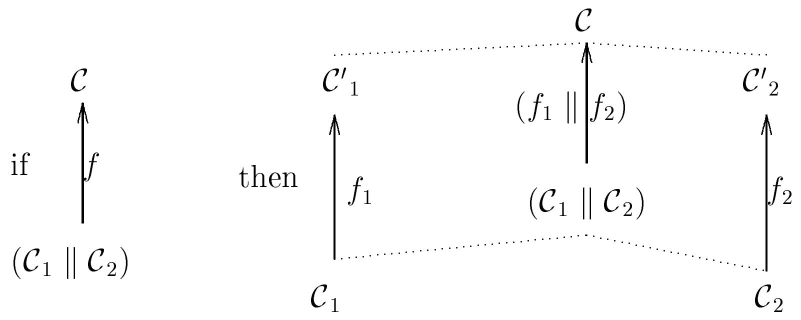

Figure 3.7: Parallel Subchannel Constraint.

3.5.4 Serial Subchannel Constraint

The Serial Subchannel Constraint is a direct analogue to the Parallel Subchannel Constraint, but requires more work. To see why, consider channels $\mathcal{C}_1, \mathcal{C}_2$, containing constraints $\phi \rightarrow \psi$ and $\psi \rightarrow \tau$ respectively. Suppose (some of) the background conditions of these constraints are made explicit in channels $\mathcal{C'}_1, \mathcal{C'}_2$ respectively. Following the assumption that such conditions are contained in antecedents, this requires the existence of constraints $\phi' \rightarrow \psi$ and $\psi' \rightarrow \tau$, where $f(\phi \rightarrow \psi) = \phi' \rightarrow \psi$ and $f(\psi \rightarrow \tau) = \psi' \rightarrow \tau$. Now, even though $\mathcal{C}_1$ and $\mathcal{C}_2$ can be serially composed, $\mathcal{C'}_1$ and $\mathcal{C'}_2$ cannot, since ψ and ψ' are (presumably) distinct types. However, the Consequent Background Constraint ensures the presence of another channel $\mathcal{C''}_2$ in the hierarchy, containing the constraint $\psi \rightarrow \tau''$, such that $\mathcal{C}_2 \sqsubseteq_{f''} \mathcal{C''}_2$ and $f(\psi \rightarrow \tau) = \psi \rightarrow \tau''$.

I can now provide definitions and results exactly analogously to those of the previous section.[49]

[49] The results are *almost* exactly analogous—the consequent of the constraint of the left-hand channel has to match the antecedent of the right-hand channel (this condition did not need to be checked for Parallel Subchannel).

Definition *Let* $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C'}_1, \mathcal{C'}_2$ *be channels, with* $\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C'}_1$ *and* $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C'}_2$*. Further suppose* $\phi{\rightarrow}\psi \in typ(\mathcal{C}_1), \psi{\rightarrow}\tau \in typ(\mathcal{C}_2)$*,* $\phi{\rightarrow}\tau \in typ(\mathcal{C}_1 \parallel \mathcal{C}_2)$ *and* $succ(f_1(\phi{\rightarrow}\psi)) = ante(f_2(\psi{\rightarrow}\tau))$*.*[50] *We define* $(f_1\,;\,f_2) : typ(\mathcal{C}_1\,;\,\mathcal{C}_2) \rightarrow typ(\mathcal{C'}_1\,;\,\mathcal{C'}_2)$ *to be the function that maps* $\phi{\rightarrow}\tau$ *to* $\phi^{f_1}{\rightarrow}\tau^{f_2}$*, where* $\phi^{f_1} = ante(f_1(\phi{\rightarrow}\psi))$ *and* $\tau^{f_2} = succ(f_2(\psi{\rightarrow}\tau))$*.*

Furthermore, the following result can be shown to hold.

Proposition 3.5.2 *Suppose all the conditions of the above definition hold. Then* $(\mathcal{C}_1\,;\,\mathcal{C}_2) \sqsubseteq_{(f_1\,;\,f_2)} (\mathcal{C'}_1\,;\,\mathcal{C'}_2)$*.*

Finally, the Serial Subchannel Constraint asserts the converse of the property proved immediately above.

Definition Serial Subchannel Constraint: *If* $(\mathcal{C}_1\,;\,\mathcal{C}_2) \sqsubseteq_f \mathcal{C}$*, then there exist* $\mathcal{C'}_1, \mathcal{C'}_2$*, where* $\mathcal{C}_1 \sqsubseteq_{f_1} \mathcal{C'}_1$ *and* $\mathcal{C}_2 \sqsubseteq_{f_2} \mathcal{C'}_2$*, such that* $\mathcal{C} = (\mathcal{C'}_1\,;\,\mathcal{C'}_2)$ *and* $f = (f_1\,;\,f_2)$*.*

This constraint is illustrated in Figure 3.8.

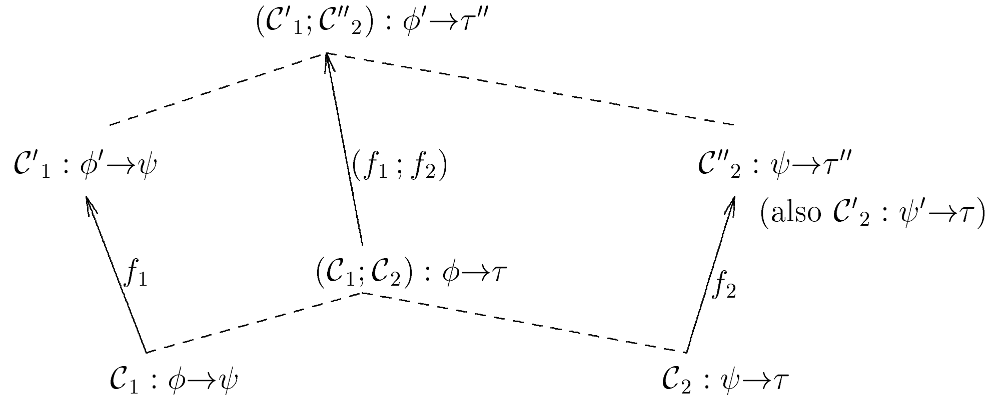

Figure 3.8: Example using background assumptions in consequent for serial composition.

The constraints imposed on the channel hierarchy so far ensure the existence of sufficient channels that (i) validate the assumption that background conditions are represented in antecedents, and (ii) smooth out the interaction between the channel operations and the channel hierarchy. The rest of the constraints are directed towards

[50] As I noted, I need to check this last condition. Each of $\mathcal{C}_1$ and $\mathcal{C}_2$ may have several superchannels, and it is only those which satisfy this condition, allowing them to be serially composed, to which this definition is applicable.

ensuring that the channel hierarchy is "internally consistent", in a way that is discussed below. Even though these constraints are independently motivated for this reason, they also serve to ensure that certain desirable patterns of reasoning are supported by the channel theoretic conditional logic.

3.5.5 Reliability Constraint

The implicit background assumptions of a regularity, supported by a channel $\mathcal{C}$, are represented in superchannels of $\mathcal{C}$. These assumptions may be distributed throughout the channel hierarchy—i.e. different assumptions are made explicit in different superchannels. However, no such assumptions should mutually conflict. This criterion is captured by the following constraint.

Definition Reliability Constraint: *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel containing types* $\phi{\rightarrow}\psi$ *and* $\phi{\rightarrow}\tau$. *Let* $\mathcal{C}_1, \mathcal{C}_2$ *be such that* $\mathcal{C} \sqsubseteq_{f_1} \mathcal{C}_1$ *and* $\mathcal{C} \sqsubseteq_{f_2} \mathcal{C}_2$. *Then it is not the case that* $ante(f_1(\phi{\rightarrow}\psi)) \perp ante(f_2(\phi{\rightarrow}\tau))$.

The Reliability Constraint is illustrated in Figure 3.9.

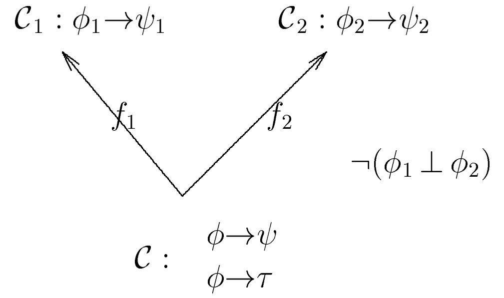

Figure 3.9: Reliability Constraint.

If a channel contains two constraints with the same antecedent, the Reliability Constraint ensures that these constraints have mutually consistent background assumptions. If this were not the case, then the channel $\mathcal{C}$ would be inherently "unreliable", or "unrobust", that it would be *guaranteed* to contain exceptions—i.e. any connection that satisfied the antecedent ϕ would necessarily be an exception to one constraint or the other.[51]

[51] This is assuming that classifications are coherent, as has been assumed throughout.

Proposition 3.5.3 *Suppose in the previous definition that*

$$ante(f_1(\phi \rightarrow \psi)) \;\perp\; ante(f_2(\phi \rightarrow \tau))$$

and let $(a \mapsto b) \in tok(\mathcal{C})$ *be such that* $(a : \phi)$ *holds in* A. *Then* $a \mapsto b$ *is an exception to at least one of the given constraints.*

Of course, two constraints with identical antecedents can still have conflicting background assumptions; however, the given constraints must then be contained in separate channels. Consider the following example, suggested to me by Jerry Seligman (private communication):

> "If there is sugar in the coffee (and it is daytime) then it tastes good."
> "If there is sugar in the coffee (and it is nighttime) then it tastes good."

The parenthesised part of each conditional is meant to be taken as an implicit background assumption. Both these conditionals (when implicit assumptions are ignored) are represented via a constraint of the form $sugar \rightarrow good$. However, they are different regularities—one concerns the goodness of sweet coffee during the day, while the other concerns the goodness of sweet coffee at night. In particular, the two conditionals taken together are not intended to somehow assert that sweet coffee always tastes good. If whether or not it is daytime or nighttime is an implicit *condition* behind the regularity, then there *must* be two separate regularities involved, even though they link the same types.

The Reliability Constraint proves extremely useful below, ensuring certain desirable patterns of inference. However, the proposition above, concerning the "unreliability" of channels that fail to satisfy this constraint, shows that Reliability is independently motivated and should be imposed on *any* channel hierarchy, regardless of the resulting inference pattern that arises. In fact, an even stronger version of Reliability can be asserted, as follows:

Definition Reliability Constraint (Revised): *Let* $\mathcal{C} : A \Rightarrow B$ *be a channel containing types* $\phi_1 \rightarrow \psi_1, ..., \phi_n \rightarrow \psi_n$, *with, for each* i, j, *either* $\phi_i \leq_A \phi_j$ *or* $\phi_j \leq_A \phi_i$.[52] *Let* $\mathcal{C}_1, ..., \mathcal{C}_n$ *be such that* $\mathcal{C} \sqsubseteq_{f_i} \mathcal{C}_i$ *for each* i. *Then all the* $ante(f_i(\phi_i \rightarrow \psi_i))$ *are mutually consistent.*[53]

It can be shown that violation of the revised Reliability constraint leads to a lack

[52] I.e. each pair of the antecedents is comparable in the type-ordering associated with A.
[53] I.e. $\bigwedge_{i=1}^{n} ante(f_i(\phi_i \rightarrow \psi_i))$ is a consistent type.

of robustness, just as with the simpler version.[54] When I make use of the Reliability Constraint below, I will mean this revised, more powerful version.

3.5.6 Consequent Consistency Constraint

The final constraint I define here is another that relates to the internal consistency of channels within a hierarchy. This particular constraint is only meaningful within the context of the assumption that all signalling relations are "reflexive"—i.e. the source and target of each connection are the same token.

The Consequent Consistency Constraint is straightforward—it ensures that background conditions associated with the antecedent of a constraint do not conflict with the consequent of that constraint.

Definition Consequent Consistency Condition: *Let $\mathcal{C} : A \Rightarrow A$ be a channel with $\phi \to \psi \in typ(\mathcal{C})$. Then for any channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, it is not the case that $ante(f(\phi \to \psi)) \perp \psi$.*

Not surprisingly, any channel that failed to satisfy this constraint would fail to be robust in a similar manner to one that invalidates the Reliability Constraint.

Proposition 3.5.4 *Let $\mathcal{C}$ be a channel as described in the above definition, and $\phi \to \psi \in typ(\mathcal{C})$ a constraint that fails the given condition. If $a \mapsto a \in tok(\mathcal{C})$ and $(a : \phi)$ holds in A, then $a \mapsto a$ is an exception to $\phi \to \psi$.*

3.5.7 Summary

In this section, I have defined a number of conditions that all channel hierarchies are henceforth assumed to satisfy:

- Antecedent Background Constraint;
- Consequent Background Constraint;
- Parallel Subchannel Constraint;
- Serial Subchannel Constraint;

---

[54] Note that the previous Reliability condition is a special case of the revised condition—it corresponds to the case where $n = 2$ and $\phi_1 = \phi_2$.

- Reliability Constraint;
- Consequent Consistency Constraint.

Of these, the first two simply ensure that there are sufficient channels in the hierarchy so that any background assumption is represented in the antecedent (respectively, consequent) of a superchannel; the next two smooth out the interaction between the channel hierarchy and the composition operations; while the final two simply ensure the internal coherence of individual channels within a hierarchy. In particular, there are no constraints whose sole purpose is to provide specific patterns of inference. However, as is shown in Section 3.6, these few constraints on a hierarchy ensure many of the patterns of inference desired of conditional logics.[55]

Note that the above constraints do not fully define a channel hierarchy—there are many hierarchies that will fit the required constraints. This is part of the relativism of the channel theoretic model—different hierarchies result in different logics of conditionals. However, any candidate hierarchy is assumed to satisfy these constraints, resulting in a "minimum" guaranteed behaviour.

3.6 Patterns of Inference

The notion of a channel hierarchy was introduced above as a mechanism for encoding implicit background assumptions of a channel, allowing the channel operations to be re-defined in a way that invalidate patterns of inference such as Transitivity, Strengthening of Antecedents, and Contraposition. However, many other axioms and patterns of inference seem valid for conditionals. The question that I turn to here is whether such patterns of inference are supported by the channel-theoretic model of conditionals presented above.

3.6.1 Hierarchy as a Parameter or Variable

I briefly mentioned above that a particular hierarchy of channels could be seen as another parameter in a relativistic model of conditionals—i.e. whether a conditional sentence is

[55] Of course, further constraints could be added, or the ones above modified, to obtain other patterns of inference, if desired.

assertible or not can only be judged within the context of a specific hierarchy of channels. An alternative view is to take a more model-theoretic approach and treat a hierarchy as a *variable*—an inference is considered valid if it holds across *all* possible hierarchies. My preferred view is the former—an agent will accept a particular conditional, or a given pattern of reasoning involving various conditionals, if it is licensed by certain aspects of his or her internal and external state, aspects of which are modelled by a particular channel hierarchy. However, there are some patterns of reasoning that should be supported by *any* channel hierarchy—these will correspond to the acceptable axioms and rules of inference of a traditional conditional logic. These patterns are examined in the following section.

Relativising inference with respect to a given channel hierarchy introduces the possibility of seemingly reasonable conclusions being invalidated because of some aspect of the structure of the hierarchy in question. For example, consider the usual doorbell example, but with the associated channel hierarchy as (partially) given in Figure 3.10.[56] This situation is unusual in that $\neg ringing$ is inserted as a background condition to $\mathcal{C}_2 : pushed \rightarrow occupied$.

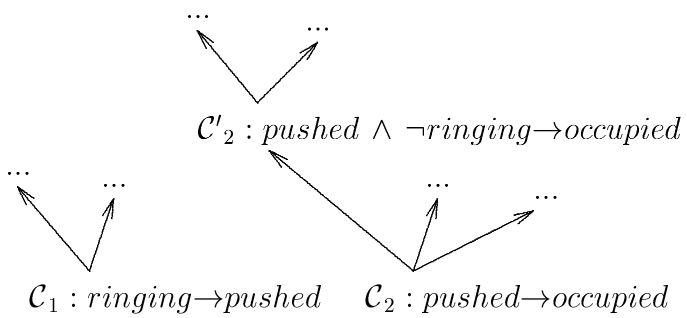

Figure 3.10: Doorbell Example with Strange Background Condition.

By the definition of the conditional serial composition operation, the composition channel $(\mathcal{C}_1 ; \mathcal{C}_2)$ does not contain the constraint $ringing \rightarrow occupied$ since the antecedent of the constraint in $\mathcal{C}'_2$ conflicts with $ringing$. At first, this seems problematic—the connection between a ringing doorbell and a pushed button seems to be independent of the connection between a pushed button and an occupied porch and so the composition

[56]This example was suggested to me by Peter Ruhrberg.

should be allowed in all cases. However, the hierarchy illustrated in Figure 3.10 does explicitly include such a dependency—the type $\neg ringing$ is an implicit background condition to the constraint *pushed→occupied*. This means that (within the context modelled by the given hierarchy) this constraint is only reliable in channel $\mathcal{C}_2$ when *ringing* does not hold of the appropriate token.[57] In this light, the behaviour of the channel theoretic logic seems perfectly reasonable for this example.

3.6.2 Axioms of Conditional Logic

Ever since Stalnaker first proposed the controversial *Conditional Excluded Middle (CEM)* as an axiom of his conditional logic, there has been disagreement as to which axioms are acceptable for conditional logics. Lewis' objection to Stalnaker's Uniqueness Assumption (which leads to CEM) led him to develop his *System of Spheres* model theoretic semantics. This semantics is sound and complete with respect to a logic that invalidates CEM, but contains other axioms which I consider objectionable. Nute (1980, 1984) reviews several important systems of conditional logic and critically evaluates the acceptability of the axioms of the various logics. The rules and axioms considered by Nute are summarised in Table 3.6.2.[58] In the table, the symbols are interpreted as follows: $\rightarrow$ is the conditional connective; $\supset$ is material implication; $\Leftrightarrow$ is logical equivalence (with respect to material implication—i.e. $A \Leftrightarrow B =_{def} (A \supset B \wedge B \supset A)$).

[57] The appropriate token is the bell connected to the button by some connection in $\mathcal{C}_1$. This is, of course, a little vague—a more rigorous formal treatment requires dropping the simplifying assumption that all channels are reflexive.

[58] I have omitted axioms and rules concerned with modalities. For simplicity, I have also omitted Nute's rules of substitution.

	Rules
RCE	if $A \supset B$ then $A \rightarrow B$
RCEA	if $A \Leftrightarrow B$ then $(C \rightarrow A) \Leftrightarrow (C \rightarrow B)$
RCEC	if $A \Leftrightarrow B$ then $(A \rightarrow C) \Leftrightarrow (B \rightarrow C)$
RCK	if $(A_1 \wedge \ldots \wedge A_n) \supset B$ then $((C \rightarrow A_1) \wedge \ldots \wedge (C \rightarrow A_n)) \supset (C \rightarrow B)$
	Axioms
A3	$((A \rightarrow (B \supset C)) \wedge (A \rightarrow B)) \supset (A \rightarrow C)$
MP	$(A \rightarrow B) \supset (A \supset B)$
RT	$((A \wedge B) \rightarrow C) \supset ((A \rightarrow B) \supset (A \rightarrow C))$
CA	$((A \rightarrow B) \wedge (C \rightarrow B)) \supset ((A \vee C) \rightarrow B)$
CC	$((A \rightarrow B) \wedge (A \rightarrow C)) \supset (A \rightarrow (B \wedge C))$
CEM	$(A \rightarrow B) \vee (A \rightarrow \neg B)$
CG	$((A \rightarrow B) \wedge (A \rightarrow \neg B)) \Leftrightarrow (A \rightarrow (B \wedge \neg B))$
CM	$(A \rightarrow (B \wedge C)) \supset ((A \rightarrow B) \wedge (A \rightarrow C))$
CS	$(A \wedge B) \supset (A \rightarrow B)$
CSO	$((A \rightarrow B) \wedge (B \rightarrow A)) \supset ((A \rightarrow C) \Leftrightarrow (B \rightarrow C))$
CV	$((A \rightarrow B) \wedge \neg(A \rightarrow \neg C)) \supset ((A \wedge C) \rightarrow B)$
ID	$A \rightarrow A$
MOD	$(\neg A \rightarrow A) \supset (B \rightarrow A)$
S*	$(\neg(A \wedge B) \rightarrow C) \supset ((\neg A \rightarrow C) \wedge (\neg B \rightarrow C))$
SDA	$((A \vee B) \rightarrow C) \supset ((A \rightarrow C) \wedge (B \rightarrow C))$
ST10	$(A \rightarrow B) \Leftrightarrow (\neg\neg A \rightarrow B)$

Table 3.6.2: Axioms and Rules of Inference discussed by Nute.

Nute critically examines arguments for and against the above rules and axioms. Some of these axioms are (at least historically) of greater interest than others. *CEM* is the characteristic axiom of Stalnaker's logic, and follows from his uniqueness assumption (i.e. whereby each world has a unique "nearest" world). *MOD, CS* and *CV* all hold in Lewis' logic (as well as Stalnaker's). Of these, *MOD* and *CS* are clearly unsuitable for a logic of necessitation conditionals (i.e. conditionals that require some sort of connection between antecedent and consequent). While many authors accept *CS* as a valid axiom for simple conditionals (i.e. those not requiring a connection), Nute presents arguments for rejecting *CS* altogether (e.g. (Nute 1980, p. 68)).

The axiom *CV* is another axiom that seems attractive, encoding a weakened form of Monotonicity. This axiom seems particularly attractive from the viewpoint of using conditional logic as a basis for default reasoning. However, Nute (following Pollock (1976)) also rejects this axiom (Nute 1980, p. 69). Similarly, the axiom *RT* encodes a weakened form of Transitivity.[59] This axiom is accepted by Nute, and is one of the axioms in the conditional logic that forms the basis of Delgrande's (1988) logic of default reasoning.[60]

Of the rules of inference, Nute rejects *RCEA,* mainly on the grounds that any logic that contains both *RCEA* and either of *SDA* or *S** must also support Monotonicity. This commits Nute to a *non-classical* conditional logic.[61] Nute devotes much effort to arguing the case for non-classical logics of conditionals, based mainly on evidence supporting the acceptability of axioms *SDA* and *S*.* Since possible-worlds semantics that treat propositions as sets of worlds cannot distinguish logically equivalent propositions, standard possible-worlds treatments of conditionals are classical.[62] Nute finds all the other rules of inference in Table 3.6.2 acceptable.

The conditional logics of Mares and Fuhrmann (1993) and of Hunter (1980) have in common the property that they are based on relevant logics.[63] I see their approaches as being related to the channel theoretic model in that they are concerned with informational links between antecedent and consequent[64] and do not support axioms such as *MOD* and *CS.* Mares and Fuhrmann's logic **ConR** and Hunter's logic of indicatives are both (relevantly) classical, in the sense that they support the relevant versions of the rules *RCEC* and *RCEA*. This is, of course, a much weaker criterion than standard *RCEC* and *RCEA* and Nute's (1980) simple proof that *RCEA* and *SDA* together lead to Monotonicity is actually invalidated when the implication in *RCEA* is relevant implication. While it is unclear whether *RCEA* and *SDA* can be simultaneously accommodated in a logic of conditionals that does not support Monotonicity, Hunter presents a case for the rejection of *SDA* altogether using an example of the following schematic form:

[59]This is perhaps more easily seen if *RT* is written as: $((A \to B) \wedge ((A \wedge B) \to C)) \supset (B \to C)$.

[60]Delgrande's (1988) logic also happens to contain *CV.*

[61]Following Chellas (1975), Nute defines a conditional logic to be *classical* if it supports the rules *RCEA* and *RCEC.*

[62]Nute uses a non-standard definition of substitution to avoid this problem.

[63]Mares and Fuhrmann present both a Hilbert-style axiomatisation and a Routley-Meyer semantics for their logic. Hunter provides only Hilbert-style axiomatisations, one for a logic of indicatives and one for a logic of subjunctives (for simplicity, I will limit the following discussion to the former of these).

[64]As noted earlier, logicians working on relevant logic claim it to be a logic of information.

$$((A \vee \neg A) \to A) \supset ((A \to A) \wedge (A \to \neg A)).$$

More controversially, Hunter also supports Contraposition, which thereby requires him to reject *CM,* since he notes that both *CM* and *SDA* are inter-derivable in the presence of Contraposition. **ConR**, on the other hand, contains *CM* (and does not support Contraposition).

I will return briefly to the adequacy of specific patterns of inference for conditionals when I evaluate the channel-theoretic logic. For now, I simply note that the validity of such patterns has been disputed in the philosophical literature over many years, with significant disagreements and changes of generally accepted rules and axioms. The brief summary of this section, however, provides a basis under which the patterns of inference supported by the channel-theoretic conditional logic can be discussed.

3.6.3 Comparing the Channel-Theoretic Logic to Standard Conditional Logics

Interpreting Standard Axioms and Rules

The first issue to address in the evaluation of the channel theoretic model of conditional logic is how the rules and axioms of standard conditional logics should be interpreted in the channel theoretic framework. A rigorous treatment of such an interpretation is far from straightforward because of the contextual aspect introduced by a channel—i.e. different conditionals may be supported by different channels, even though those conditionals can be composed in some straightforward manner. This is complicated further since any resulting conditional will generally be supported by yet another different channel, one which somehow merges the contexts of the channels supporting the original conditionals.[65] To simplify matters, I will make the following assumptions.[66] Firstly, continuing the assumption made earlier, the token-level connections of channels are assumed to be reflexive. Without this assumption, the following comparison between the channel-theoretic and standard conditional logics would not be meaningful since the way a conditional was "situated" (i.e. related to token-level connections) would have to be taken into account. This allows me to effectively continue to ignore the token-level

[65] I.e. a channel formed by serial or parallel composition is usually different to both the channels from which it is composed.

[66] Appendix B discusses the relaxation of these assumptions.

connections of channels, circumventing possible problems related to the interpretation of negated conditionals. Secondly, I consider only "flat" axiom schemata—i.e. nested conditionals and nested use of the material implication are ignored because of the simplification made earlier (but see Appendix B.1). Finally, I consider only axioms whose consequent is a conditional (i.e. the prominent connective in the consequent is the conditional operator). This rules out axioms *SDA* and *S** from consideration—these axioms are discussed separately later.

A major difference between a possible-worlds view of conditionals and the channel theoretic model is in the treatment of context. For example, in Stalnaker's nearest-world semantics, an axiom such as CC is true in a world w if the consequent of CC (i.e. $A \rightarrow (B \wedge C)$) is true in w or if the antecedent (i.e. $(A \rightarrow B) \wedge (A \rightarrow C)$) is false in w. The truth of the conditional sentences embedded in the antecedent and consequent of CC is determined by the worlds accessible from w, but each of these conditionals is effectively evaluated in the "same" context, since they are all evaluated at the same world w under the same "nearest-world" function. In the channel theoretic model, each of the embedded conditionals of the antecedent may well be supported by different channels $\mathcal{C}_1$ and $\mathcal{C}_2$. These channels may incorporate conflicting background assumptions—for example, $\mathcal{C}_1$ may support $A \rightarrow B$ under the assumption that $\neg C$ holds, while $\mathcal{C}_2$ supports $A \rightarrow C$ under the assumption that C holds. The appropriate "merger" of these channels (i.e. the one arrived at by the parallel meet operation) should not, in such a case, support the conditional $A \rightarrow (B \wedge C)$.

Clearly, I need some concept of " the same context" in the channel theoretic model before I can meaningfully discuss which axioms from Table 3.6.2 are supported by the channel theoretic logic of conditionals. One possibility is to precisely define what it means for two channels to capture the "same", or at least "compatible", contexts. This could be done by examining the channels above them in the pertinent channel hierarchy, but I will take the simpler option of assuming that two channels provide the same context only if they are the same channel. From this it follows that, for the channel theoretic model to support an axiom Ax of conditional logic, I require that, for every channel $\mathcal{C}$ such that $\mathcal{C}$ supports each embedded conditional in the antecedent of Ax, some "appropriate" channel (see below) supports the consequent.[67]

[67] I can restrict myself to the case where $\mathcal{C}$ is required to support *all* conditionals in the antecedent of Ax since each such antecedent contains only conjunctions of conditionals.

The other question to be addressed regards the appropriate channel $\mathcal{C}'$ that supports the consequent of an axiom. Since the operations forming the basis of the channel theoretic logic of conditionals "merge" channels, $\mathcal{C}'$ need not be the same channel as that supporting the conditionals in the antecedent. I could assume that channels are *closed* under application of the operations,[68] the same way in which logical channels are, but I will instead take the view that an appropriate channel for the consequent of an axiom is one that is obtained by the use of the channel operations on $\mathcal{C}$ (the channel supporting the conditionals in the antecedent of the axiom) and any logical channel $\mathcal{L}$.[69] So, for example, the channel theoretic model supports the axiom CA from Table 3.6.2 since, for any channel $\mathcal{C}$ supporting $A \rightarrow B$ and $C \rightarrow B$, the parallel composition of $\mathcal{C}$ with itself (i.e. $(\mathcal{C} \parallel \mathcal{C})$) supports $(A \vee C) \rightarrow B$.

The above discussion can be made more precise by the following definition of what it means for a channel hierarchy to *support* an axiom of conditional logic.[70] In the following, it is assumed that there is some mapping * that takes sentences of some logical language **L**, in which the axioms of Table 3.6.2 are stated, to (structured) types—in particular, the logical symbols (conjunction, disjunction and negation) are modelled by the corresponding type-operations.[71] It is straightforward (though slightly messy) to define such a mapping. Also, due to the simplification made earlier whereby all connections are reflexive, the token-level is ignored, thereby greatly simplifying the definition.

Definition *A channel* $\mathcal{C}$ supports *a sentence* S, *written* $\mathcal{C} \models S$, *if the following holds:*

1. $\mathcal{C} \models A \wedge B$ *iff* $\mathcal{C} \models A$ *and* $\mathcal{C} \models B$*;*
2. $\mathcal{C} \models A \vee B$ *iff* $\mathcal{C} \models A$ *or* $\mathcal{C} \models B$*;*
3. $\mathcal{C} \models \neg A$ *iff it is not the case that* $\mathcal{C} \models A$*;*
4. $\mathcal{C} \models A \rightarrow B$ *iff* $A^* \rightarrow B^* \in typ(\mathcal{C})$.[72]

Given a channel hierarchy $\mathcal{H}$ *and a channel* $\mathcal{C}$ *in* $\mathcal{H}$*, let* $a(\mathcal{C})$ *(the channels "accessible" from* $\mathcal{C}$*) be the collection of channels that are obtained from* $\mathcal{C}$ *by use of the conditional*

[68] That is, I could assume that $(\mathcal{C}\,;\mathcal{C}) = \mathcal{C}$ and $(\mathcal{C} \parallel \mathcal{C}) = \mathcal{C}$.

[69] I can allow indiscriminate use of logical channels to prove axioms since these channels effectively do not participate in the channel hierarchy—i.e. they have no associated restricted context.

[70] Note that the following definition is not fully general but captures the notions discussed above. In particular, while it may seem strange to define the notion of a channel supporting a sentence such as $\neg A$, this is unproblematic since, given the sentences under consideration, A is never anything other than a conditional.

[71] E.g., $(A \wedge B)^* = A^* \wedge B^*$, where the second use of $\wedge$ is type-conjunction.

[72] As always, I have used the same symbol for the conditional connective of **L** and the channel-theoretic concept of indicating relation—hopefully this will cause no confusion.

composition operations to $\mathcal{C}$ and any logical channel $\mathcal{L}$ (as described earlier).[73] *A hierarchy $\mathcal{H}$ is said to* validate *a sentence S, written $\mathcal{H} \models S$, if the following holds:*

1. *$\mathcal{H} \models A \to B$ iff $\mathcal{L} \models A \to B$ for some logical channel $\mathcal{L}$ in $\mathcal{H}$;*
2. *if A and B do not contain any conditional connectives, then*
   $\mathcal{H} \models A \supset B$ iff $A^* \to B^* \in typ(\mathcal{L})$ for some logical channel $\mathcal{L}$;
3. *if A or B does contain a conditional connective, then*
   $\mathcal{H} \models A \supset B$ iff for every channel $\mathcal{C}$ in $\mathcal{H}$, if $\mathcal{C} \models A$ then $\mathcal{C}' \models B$ for some $\mathcal{C}' \in a(\mathcal{C})$;
4. *otherwise,*
   $\mathcal{H} \models A$ iff $\mathcal{C} \models A$ for every channel $\mathcal{C}$ in $\mathcal{H}$.

In the above definition, an axiom is supported in a hierarchy if either it is supported by a logical channel (i.e., a channel that models analytic information-flow), or it is supported by every channel in the hierarchy. The axioms of greatest interest are those of the form $A \supset B$, where A and B contain conditionals, and these are handled in the way discussed above—i.e. if $\mathcal{C} \models A$, then for some channel $\mathcal{C}'$ accessible from $\mathcal{C}$ (and $\mathcal{L}$) via the channel operations, $\mathcal{C}' \models B$. Notice that axioms such as *CM, SDA* and *S*,* i.e., axioms with a conjunction of conditionals in the consequent, are not really adequately handled by this definition since they involve conjunctions of conditionals in the consequent of the axiom. These axioms are discussed below.

The interpretation of rules of inference is straightforward given the above concept of a hierarchy validating a logical sentence—if the premise of the rule is validated by the hierarchy, then so too must be the conclusion.

Definition *A channel hierarchy $\mathcal{H}$ is said to* validate *a rule of inference if the following holds:*

> *if $\mathcal{H}$ validates the premise of the rule, then $\mathcal{H}$ also validates the conclusion.*

Evaluation

Given the interpretation of axioms and rules of standard conditional logics in the channel theoretic model, one can determine which of those contained in Table 3.6.2 are validated. Of course, exactly which patterns of inference are validated by a particular model is

[73] $a(\mathcal{C})$ could be defined as the smallest collection of channels containing $\mathcal{C}$ and $\mathcal{L}$ and closed under the conditional channel operations.

dependent on the associated channel hierarchy. However, some patterns hold *regardless* of the channel hierarchy, and it is this sort of result I am interested in. All non-trivial proofs are relegated to Appendix B.5.

Consider first some of the interesting axioms. I have already pointed out that *CA* is validated regardless of the hierarchy involved. If $\mathcal{C}$ is a channel containing constraints $\phi \rightarrow \psi$ and $\tau \rightarrow \psi$, then the channel $(\mathcal{C} \parallel \mathcal{C})$ *must* contain the constraint $(\phi \vee \tau) \rightarrow \psi$ since the definition of parallel join is unmodified in the presence of a channel hierarchy. Other axioms are just as easily seen to be invalidated in the channel theoretic interpretation. For example, *CEM, CS* and *MOD* all make claims that are not supported in the highly intensional channel theoretic model of conditionals. (This is in agreement with the relevant logic models of conditionals (Hunter 1980; Mares and Fuhrmann 1993).) Also, *ID* is supported by any logical channel.

Less obvious is the fact that any channel hierarchy satisfying the conditions of Section 3.5 validates the axiom *CC.* Given a channel $\mathcal{C}$ containing constraints $\phi \rightarrow \psi$ and $\phi \rightarrow \tau$, it can be shown, using the Reliability condition, that the channel $(\mathcal{C} \parallel \mathcal{C})$ contains $\phi \rightarrow (\psi \wedge \tau)$ (see Appendix B.5). Also, it can be shown that Nute's *ST10* is supported (again, see Appendix B.5). One axiom from Table 3.6.2 that isn't validated by the channel theoretic model is *CV.* This is a form of weakened Transitivity that is supported in many possible-worlds treatments of conditionals. However, the fact that a channel $\mathcal{C}$ fails to contain a constraint of the form $\phi \rightarrow \neg\tau$ does not prevent the possibility of some channel $\mathcal{C}'$, such that $\mathcal{C} \sqsubseteq \mathcal{C}'$, containing a constraint of the form $\phi' \rightarrow \tau$, where $\phi' \perp \tau$. Thus the serial composition of a logical channel $\mathcal{L}$, containing $(\phi \wedge \tau) \rightarrow \phi$, and $\mathcal{C}$, containing $\phi \rightarrow \psi$, need not contain the constraint $(\phi \wedge \tau) \rightarrow \psi$. As discussed above, this is in agreement with Nute (1980) and Pollock (1976), who provide various reasons for rejecting *CV.* The channel theoretic model does, however, validate Delgrande's version of *RT,* which is also a weakened form of Transitivity (see Appendix B.5). This time, Nute agrees with the validity of this axiom.

So far, the patterns of inference validated by the channel theoretic model have been very closely aligned to those argued for by Nute. However, an important deviation is provided by the axioms *SDA* and *S*.* Consider *SDA.* What is required for a channel hierarchy to validate such a pattern of inference is for every channel supporting a constraint with a disjunction in the antecedent to be able to be "decomposed"—i.e. any such channel $\mathcal{C}$ would have to be the parallel composition (specifically, the parallel join) of two other channels. Such a constraint could be imposed on channel hierarchies, in addition to those defined in Section 3.5. This would result in a rather attractive symmetry to a hierarchy—not only could any pair of channels be composed to form a channel in the hierarchy, but any channel would be decomposable to two other channels in the hierarchy. However, Hunter (1980) has argued that the pattern of inference represented by *SDA* should not be considered valid. Another axiom that would require the condition that channels be decomposable is *CM.* This is rejected by Hunter, since it is derivable from *SDA* using Contraposition, but is supported by Mares and Fuhrmann (who do not support Contraposition). While I cannot provide an independent counter-example for *CM,* I will not pursue this issue any further here, but leave it as an interesting possible way in which the channel theoretic model may be extended.

I now turn to the rules of inference in Table 3.6.2. It turns out the channel theoretic model validates both $RCEC$ and $RCEA$ making it a (relevant!) classical conditional logic. Of course, this is not a problem since *SDA* is not validated in general. However, even if it was, logical equivalence in relevant logic is weaker than in classical, and Nute's argument for rejecting $RCEA$ in the presence of *SDA* is nullified. The relevant versions of RCE and RCK are also validated. All these results are shown in Appendix B.5.

3.7 Discussion

A channel hierarchy provides a parameter against which the patterns of inference of a particular channel theoretic model of conditionals may vary—i.e. changing the hierarchy may result in different patterns of inference. The previous section was concerned with patterns of reasoning that hold with respect to *any* channel hierarchy satisfying the constraints defined in Section 3.5, such as Reliability and Consequent Consistency. Other constraints could be imposed on hierarchies, thereby guaranteeing other patterns of reasoning across all such hierarchies. One possibility was mentioned above, whereby *CM* seems to require any channel to be decomposable into constituent channels. However, the fairly simple, natural conditions[74] of Section 3.5 already ensure quite powerful patterns of reasoning for conditionals.

The channel hierarchy framework introduced in this chapter is also exploited in each

[74]By "natural", I mean that the conditions do not require more than should be required of *any* channel in general.

of the following chapters. The importance of the framework is that it allows background assumptions of a regularity to be made explicit enough so that they can be accounted for in a calculus of channels—as demonstrated in this chapter—while still allowing a channel to be used in inference *independently* of any hierarchy in which it is to be found. This is an important property if the framework is to provide an adequate model of situated reasoning. This topic is further discussed in Chapter 5.

A number of simplifications were made in this Chapter, some of which are carried over into later chapters. Ways in which some of these simplifications can be removed are discussed in Appendix B, as are some possible extensions to the basic framework described here. Chapter 6 also discusses improvements and possible modifications to the channel hierarchy framework and the general approach to obtaining a channel theoretic logic of conditionals.

Chapter 4

A Channel Theoretic Model for Reasoning with Generics

In the previous chapter, channel theory was used to provide an analysis of conditional statements. A conditional sentence was taken to be an assertion that a certain channel supported some specified conditional information. In this chapter, I use channel theory as the basis of an analysis of generic statements. Once again, my main concern is with providing an adequate logic for reasoning with generics, but a channel theoretic view of generics seems to provide some definite benefits.

Recent work in the AI and philosophical literature has been concerned with the relationship between logics of conditionals and logics of generics and default reasoning (e.g. (Boutilier 1992; Delgrande 1988; Morreau 1992a)). Similarly to conditionals, generics seem to fail to validate standard patterns of reasoning, such as Transitivity, Monotonicity and Contraposition.[1] However, the move from conditional logic to default logic is never as smooth as initial similarities suggest it may be, leading to what Morreau (1992a) refers to as *"ghosts in the machine"*—*ad hoc* additions to the machinery of the conditional logic that address some specific problem that arises in the resulting default logic.[2] The logic defined in Section 4.3 builds directly on the channel theoretic logic of

[1]As is the case with conditional logics, salient counterexamples to Contraposition are hard to come by. However, Kraus et al. (1990) have shown that for a certain class of non-monotonic consequence relations, the inclusion of Contraposition leads to Monotonicity. Kraus et al.'s results are discussed in Section 5.3.1.

[2]E.g. Delgrande's (1988) solution to the *Irrelevance problem;* prioritising predicates for circumscription (Lifschitz 1985). These are discussed below.

conditionals of the previous chapter. It requires the addition of no extra machinery to the logic itself. Instead, reasoning with generics is seen to require specific assumptions on the associated channel hierarchy: namely, that any background conditions encapsulated within it must be present in the given generics from which the inferences are drawn. This is a significant shift from modifying the logic of reasoning itself, and seems to succinctly capture an important difference between acceptable logics of conditionals and logics of generics.

Before defining the logic of generics and default rules, I briefly present a simple channel theoretic analysis of generics. The major concern of a formal semantics of generics is the possibility of *exceptions* and the problem of how to cater for them. As discussed in previous chapters, this concern is one of the central features of channel theory. A channel theoretic analysis of generics seems to offer several other benefits—these are discussed through the use of several examples that have received some attention in the literature. However, my main concern is to provide a logic of generics, based on the framework for reasoning with conditionals developed in the previous chapter. This logic will be seen to satisfy many important properties; in particular, many desirable patterns of reasoning are validated, such as those specified by Asher and Morreau (1991).

4.1 The Semantics of Generics

The analysis proposed below is based on the assertion that a generic sentence describes a regularity of some sort. This in itself is a fairly uncontroversial position to take—for example, Krifka et al. (1995) (hereafter, referred to as *TGB*) assert that "a generic sentence states a law-like regularity". A more important facet of a channel theoretic analysis is the use of a channel to provide a demonstrative content to a generic sentence, thereby providing a "context" within which the regularity can be relied upon. This seems to address several issues that are problematic for other semantic accounts of generics.[3]

The issue of greatest concern in the literature on the semantics of generics has been the analysis of regularities themselves—i.e. what is it about the regularity underlying a generic sentence that makes it a regularity (in the face of possible counter-examples)? As discussed in Chapter 2, channel theory makes no reductionist claims

[3]Some of the work described in this section was performed in conjunction with Sheila Glasbey and is reported in (Cavedon and Glasbey 1994).

regarding regularities—i.e. the theory in no way provides an analysis as to when a relationship between types constitutes a regularity—but instead attempts to explain various problematic properties of regularities by providing a structural analysis of them. As such, the analysis of generics proposed below is not concerned with the question of how the regularity underlying a generic arises, but on showing that the channel theoretic view of a regularity addresses issues that have proved problematic for previous attempts at a semantics of generics. Again, however, my main concern is with providing a channel theoretic logic of generics rather than an explicit semantics at this point.

4.1.1 Characterising Sentences and Kind-Referring NPs

Krifka (1987) distinguishes two main families of generic sentences, which he calls *I-generics* and *D-generics.* In *TGB*, this same distinction is between *characterising sentences* and *kind-referring NPs,* respectively. The former sort of generic is one that asserts a general property about individuals of a particular type,[4] while the latter asserts a property of a *kind.*[5] As is generally the case with accounts of generics that are concerned with providing a logic for reasoning with them, I will be solely concerned with I-generics, or characterising sentences. Throughout, the use of the term "generic" is used to mean this particular sort, unless explicitly stated otherwise (in which case I will tend to use *TGB*'s terminology).

Characterising sentences do not seem to be syntactically marked (at least, not in English) in any particularly distinguishing way—there are many different syntactic forms in which a characterising sentence may appear. This includes the *if ... then ...* form associated with conditional sentences—these generics are the sort that Barwise (1986) calls *general conditionals.* As such, I point out once again that the distinction I make in this thesis between "conditional sentences" and "generic sentences" is not based on any syntactic difference. The analysis of this chapter is meant to cover sentences that "generalise" (as generics do) over individuals and situations, while the analysis of the previous chapter is not intended to cover such sentences. However, the two sorts of sentences—and the patterns of inference that seem valid for each—are closely related, which is a motivating factor behind the enterprise of modelling reasoning with generics using the same framework as that for modelling reasoning with conditionals. *TGB*

[4] E.g. "Birds fly".
[5] E.g. "Dinosaurs are extinct".

describes various diagnostic tests that may be applied to sentences so as to determine whether a sentence is a generic or not.

One particular sort of generic sentence that is perhaps worth distinguishing is the class of *habituals.* This class includes sentences such as the following:

"John smokes."

"Mary smokes after dinner."

TGB defines a habitual sentence to be a generic that quantifies over situations (using the term informally). Since situations simply constitute another sort of individual, this presents no special problem for the channel theoretic treatment and I will simply consider the habituals to be a subclass of the characterising sentences in general. In particular, I will not distinguish them again and will often use a habitual as a supporting example for the general approach.

4.1.2 Some Issues in the Semantics of Generics

The semantics of generic sentences has proved to be a demon in the linguistic, philosophical and AI literature for many years. The following issues are central problems in the semantics of generics and are specifically addressed by the channel theoretic analysis.

Exceptions

There are many important problematic issues that need to be addressed by a semantics of generic sentences. The most obvious of these is the need to account for the possibility of *exceptions* to the general rule. For example, the sentence

"Birds fly."

is clearly acceptable as a generic, even though there are many non-flying birds—whole subclasses of them in fact (such as penguins). While generics seem to involve some sort of quantification over individuals, this cannot be standard universal quantification. In fact, examples such as the following suggest that there is *no* simple quantifier (such as *most*) that is appropriate:

"Turtles live a long life."

This generic seems acceptable even though many (in fact, nearly all) turtles are eaten by predators while very young. Morreau (1992a) claims that the following generic would be true

"Potatoes contain vitamin C."

even if all the potatoes in the world were boiled for so long that none of them actually contained any vitamin C at all.

Various proposals for handling exceptions to general rules have appeared in both the philosophical and AI literature. In particular, there is a large body of work on *nonmonotonic* or *default reasoning* which is relevant to the semantics of generics. A particular approach to default reasoning which has recently gained popularity is the use of modal logics of *normality* (the theses of Boutilier (1992) and Morreau (1992a) are notable works in this area). Such logics involve quantification over "normal" individuals rather than *all* individuals, and thereby allow a sentence to be validated without requiring all (or, in fact, any) relevant individuals to satisfy the property claimed for them by the generic. One particularly interesting feature of these logics is that they have their basis in conditional logic, which is in common with the channel theoretic model presented below. As noted earlier, conditionals and generics seem to have many properties in common, especially if we focus on valid patterns of inference. It is therefore to be hoped that a logic for reasoning with generics could be extracted from a logic for conditionals. Boutilier's and Morreau's modal logics of normality are discussed in some detail in Section 4.1.4. The approach whereby generics are seen as quantifying over "normal" individuals is often referred to as the *normative* approach. I will use this terminology in the following.

The Intensionality of Generics

The discussion of exceptions contained examples that demonstrate that the truth of a generic is at best weakly related to properties of the individuals with which it is concerned. In the channel theoretic treatment of regularities, a regularity is itself a primitive type and the question of whether or not it is supported in a channel is independent of the behaviour of token-level connections related to it. Accounts of generics that involve quantifying over individuals can run into what Morreau (1992a) calls the *problem of extensionality.* For example, consider the following sentence (taken from (Morreau 1992a)):

"Elderly members of the club drink for free."

If the club in question does not have any elderly members, then this should not constitute a sufficient condition to validate this generic—there needs to be some sort of policy whereby the club's committee has decreed that any elderly members of the club will not be charged for their drinks. However, an account of generics that simply quantifies over some (possibly restricted) set of individuals would validate a generic such as the one above if that set of individuals was empty. In particular, Morreau (1992a) points out that accounts of generics based on *circumscription* (McCarthy 1980; McCarthy 1986) suffer from this problem.

A potential problem with taking the fine-grained approach whereby regularities are primitive objects concerns inferencing with such objects. An adequate logic of default reasoning must allow (default) conclusions to be drawn about individuals. For example, if I know that birds fly and am told that Tweety is a bird, then (in the absence of any information to the contrary) I should be able to infer that Tweety can fly.[6] In accounts of generics based on some sort of quantification (such as circumscription), the properties of (normal) individuals are a direct consequence of the truth of the pertinent generic, and vice versa. For more intensional accounts, however, an extra assumption is required, which Morreau (1992a) calls *maximal normality*—i.e. an assumption that an individual *is* in some sense "normal" needs to be made before any information regarding that individual can be inferred via the generic in question. In the channel theoretic model, this sort of assumption characterises the move from a logic of generics to a logic of default reasoning—i.e. there is no need for an assumption of normality to obtain a logic for reasoning with generics, but such an assumption is required in Chapter 5 when I extend the logic of generics to a system that is able to use generics to draw default conclusions about individuals.

The Role of Context

The fact that a generic can be seen as a statement as to what is "normally" the case can be taken as a restriction to a particular context (i.e. restricting quantification to "normal" individuals). However, other contextual factors also seem to play a role in the evaluation of generics, and these are not explained away so easily by the normative

[6] Of course, I may have to withdraw this conclusion if I am later told that Tweety is a penguin.

account of generics. Consider the following example:

"Chickens lay eggs."

Of course, not all normal chickens lay eggs, only female ones—in fact, only female ones of a certain age. Accounts of generics that involve quantifying over normal individuals generally ascribe this problem to the issue of pragmatics—for example, Morreau (1992a) claims that this issue should be resolved when deriving the correct logical form of the above sentence. However, another approach is to view the associated generic as only being pertinent to a restricted collection of individuals, in this case, female chickens of a certain age. The channel theoretic analysis described below allows this sort of view.

It also seems that different, even conflicting, generics can be acceptable in different contexts. For example, consider the following:

"Swans are white."

In various parts of the world, this is an acceptable statement. However, in other parts, notably Australia, the following statement seems acceptable:[7]

"Swans are black."

In the modal logics of normality described in Section 4.1.4, this can be handled by the fact that the ranking of worlds (with respect to normality) is not an objective matter but is defined with respect to a world. As such, with respect to world w, normal swans may be white, whereas with respect to w', normal swans are black. Proponents of this view (such as Morreau (1992a) and Boutilier (1992)) take pains to point out that there is no global sense of "normality": different agents (in different worlds) can have different views as to what is "more normal".

However, conflicting sentences can sometimes be asserted in the *same* context. For example, Cavedon and Glasbey (1994) note that a sentence such as the following:

"Mary smokes after dinner."

(with the appropriate reading being the one that allows the inference that it is after dinner from the information that Mary is smoking) can be followed by

"She smokes before breakfast, too."

[7]Certainly, (a suitable translation of) this statement seems acceptable to the people who lived in Australia before the arrival of Europeans.

(with the appropriate reading being the one analogous to that for the previous sentence) without any apparent contradiction. In the standard normative account of generics, the normal *Mary-smoking* situations cannot be both before breakfast *and* after dinner. The fact that a single agent could consistently assert (simultaneously) both the above habituals suggests that the habitual itself has a context associated with it, and that this context plays a role in its interpretation. In fact, we would expect that a "merging" of these contexts would support a sentence such as:

> "Mary smokes before breakfast and after dinner."

(This is to be interpreted roughly as: "If Mary is smoking, then it is either before breakfast or after dinner".) The channel theoretic analysis of generics presented below ascribes a demonstrative content to a generic statement, providing a coherent interpretation of the above pair of habituals.

Generics and Truth

There are two popular views of generics:

- a generic asserts a proposition that can be either true or false;
- a generic acts as a rule of (default) inference.

In the literature dealing with the formal semantics of generics and default rules, the first view is taken by proponents of circumscription (McCarthy 1980; McCarthy 1986) and modal logics of normality (Boutilier 1992; Morreau 1992a), while the second is taken by proponents of Default Logic (Reiter 1980) and Update Semantics (Veltman 1993).

A problem with the second view—generics as inference rules—is that it becomes difficult to meaningfully express the generics themselves in the logical language. In particular, none of the analyses based on that approach allows the representation of *nested* generics (i.e. whereby a generic is embedded within the premise or conclusion of some other generic). However, as argued by Morreau (1992a), there seem to be sentences which are naturally interpreted as nested generics; for example:[8]

> "Current medical theories predict that smoking leads to cancer."

[8]This example was suggested to me by Carl Vogel.

This shortcoming—the inability to represent nested generics—seems to be a serious flaw for the inference-rule interpretation of generics. Morreau (1992b) discusses the divide between the two viewpoints further, in particular showing that the valid patterns of reasoning supported under the two different approaches cannot be reconciled.[9]

Assuming that generics have truth values, the question arises as to what it is that makes a generic true. The simplest approach is to take the view that a generic is *simply true* or *simply false*—i.e. true or false "not in virtue of the properties of the individuals of those kinds, but somehow true [or false] all by themselves" (Morreau 1992a, pg. 49). Morreau's problem with this approach is that it does not seem to lead to a method for inferring default properties about individuals—e.g. if "Birds fly" is simply true, then it does not seem possible to use this to (defeasibly) infer that a given bird flies. For example, the channel theoretic logic of generics requires the definition of a "maximal normality" condition in Chapter 5 before it can be used for default reasoning (about individuals).[10]

Patterns of Reasoning

As mentioned above, some classically-valid rules of inference—such as Transitivity, Monotonicity and Contraposition—are well-accepted to be invalid for generics and default-rules. In fact, there is a fairly close correlation between the rules that are invalidated for conditional logics and those that are invalidated for logics of generics and defaults, resulting in recent attempts to define logics of default reasoning with a conditional logic basis. An acceptable logic for reasoning with generics must invalidate these undesirable patterns of inference.

The complementary issue concerns patterns of inference that do seem valid for generics. A sample list is presented by Asher and Morreau (1991). More generally, Morreau (1992a) discusses various properties that a logic of generics should satisfy. These, along with the relevant patterns of reasoning required by Asher and Morreau,[11] are sum-

[9] Koons (1993) shows that Morreau's results in this regard do not hold in the absence of Substitution of Logical Equivalents, the rule whereby logically equivalent sentences can be substituted for each other without affecting the truth of the embedding sentence.

[10] One could take the view that a clean separation between the truth of generics and their use in default reasoning is not a bad thing, however.

[11] Note that some of Asher and Morreau's patterns of reasoning are relevant to a logic of *default reasoning,* i.e. for reasoning about the (default) properties of individuals. While this is related to a logic for reasoning with generics, it requires the added mechanism of an assumption of maximal normality,

marised in the following.

- *Addition of Generics:* This is exemplified by the following example, taken from Morreau (1992a):

  "Lions are brown."
  "Lions are dangerous."

  "Lions are brown and dangerous."

  Morreau claims that this particular pattern of inference is not handled well by circumscription.

- *Graded Normality:* Carlson (1977) claims that normality should come "in degrees"—an individual being "abnormal" in one respect should not necessarily lead to that individual being considered abnormal in other respects. For example, given the two sentences in the premise of the above example, if *leo* is an albino lion (i.e. *Leo* is an abnormal lion with respect to the first generic), then *Leo* should still be considered potentially dangerous! Carlson expected Graded Normality to pose problems for normative accounts of generics, as is the case for Morreau's logic of generics (discussed below).

- *Specificity:* Also termed the *Penguin Principle* by Morreau, Specificity gained popularity in the AI literature regarding *default inheritance networks* (e.g. (Horty et al. 1990; Touretzky 1986; Touretzky et al. 1987)). It involves the claim that an inference involving two transitive generics can be overruled by a (more specific) generic that contradicts the transitive conclusion. This is exemplified by the following example:[12]

  "Penguins are birds."
  "Birds fly."
  "Penguins don't fly."

  Using Transitivity would allow the inference that "Penguins fly". However, the specific assertion of the third generic seems to prohibit such a conclusion.

  Specificity has proved problematic for all the major classical systems of default reasoning from the AI literature (e.g. Default Logic (Reiter 1980), Circumscription

---

which is defined in the following chapter. At this point, the only patterns of inference that I consider are those that are concerned solely with the generics, or default rules, themselves.

[12] This example demonstrates the invalidity of Transitivity for generics.

(McCarthy 1980) and Autoepistemic logic (Moore 1985)), leading to extensions of these frameworks whereby default rules are effectively ordered with respect to their applicability, simply in order to resolve the problem of specificity (e.g. hierarchical default logic (Touretzky 1984), prioritised circumscription (Lifschitz 1985) and hierarchical autoepistemic logic (Konolige 1989)).[13] As such, specificity presents an important testbed on which to evaluate an adequate logic of generics.

- *Irrelevant conditions:* The problem of "irrelevant conditions" is a particularly thorny one for default logics derived from a conditional logic base. Standard conditional logics are extremely weak with respect to Monotonicity, invalidating such inferences as the following:

  "Birds fly."

  "Red birds fly."

  In the absence of any link between a bird's colour and its ability to fly, a logic for defaults should be able to infer the second generic from the first. Delgrande (1988) and Boutilier (1992) each need to add extra mechanisms to their logics of conditionals before obtaining a logic of defaults that does not suffer from this problem. Morreau (1992a) sees this as a serious shortcoming of their logics.

- *Other patterns of reasoning:* The following patterns of reasoning are taken from Asher and Morreau (1991). They are those that are solely concerned with inference involving generic rules, rather than those involving the properties of individuals. (The symbol '$\rightarrow$' denotes the generic operator; the symbol '$\mid\!\approx$' denotes default consequence).

  Defeasible Transitivity:

  - $A \rightarrow B, B \rightarrow C \mid\!\approx A \rightarrow C$, *but*
  - $A \rightarrow B, B \rightarrow C, A \rightarrow \neg C \not\mid\!\approx A \rightarrow C$

  Defeasible Monotonicity:

  - $A \rightarrow B \mid\!\approx (A \wedge C) \rightarrow B$, *but*
  - $A \rightarrow B, (A \wedge C) \rightarrow \neg B \not\mid\!\approx (A \wedge C) \rightarrow B$

[13] These modifications exemplify those attacked by Morreau (1992a) as being "ghosts in the machine".

Graded Normality and Defeasible Transitivity:

- $A{\rightarrow}B, B{\rightarrow}C, B{\rightarrow}D, A{\rightarrow}\neg D \mathrel{|\!\!\leadsto} A{\rightarrow}C$, *but*
- $A{\rightarrow}B, B{\rightarrow}C, B{\rightarrow}D, A{\rightarrow}\neg D \not\mathrel{|\!\!\leadsto} A{\rightarrow}D$

These and other general patterns of inference are investigated below, after the definition of the channel theoretic framework for reasoning with generics.

4.1.3 A Dyadic Generic Operator

A very useful piece of notation for describing the semantics of generic sentences is the dyadic generic operator usually written *GEN*.[14] The *GEN* operator was introduced because of problems with previous monadic operators, the main issue being that multiple generic readings of some sentences could not be provided for using the monadic operator.[15] The general form of the operator is the following:

$$GEN[x_1, ..., x_n; y_1, ..., y_m](\Phi(x_1, ..., x_n)\ ;\ \Psi(\{x_1\}, ..., \{x_n\}, y_1, ..., y_m)).$$

Written in this form, Φ is known as the *restrictor* and Ψ as the *matrix*[16]. The notation $\Phi(x_1, ..., x_n)$ means that the x_i occur free in Φ, and the notation $\Psi(\{x_1\}, ..., \{x_n\})$ means that the x_i possibly occur free in Ψ. The $x_1, ..., x_n$ are variables bound by the *GEN* operator and $y_1, ..., y_m$ are variables that are bound existentially within Ψ, so the above can be equivalently written as follows:

$$GEN[x_1, ..., x_n;](\Phi(x_1, ..., x_n)\ ;\ \exists y_1 ... y_m \Psi(\{x_1\}, ..., \{x_n\}, y_1, ..., y_m)).$$

As an example of the use of *GEN*, consider the following habitual:

"Mary smokes after dinner."

This sentence is generally accepted to have (at least) two readings, one allowing the inference of the time from the knowledge that Mary is smoking (i.e. answering the question "When does Mary smoke?"), the second allowing an inference regarding what Mary does after dinner (i.e. answering the question "What does Mary do after dinner"?). In the standard notation using *GEN*, these two readings are written as follows, respectively:

[14] For example, see *TGB* (Section 2.3), although a number of authors have proposed similar concepts.

[15] Recent work by Krifka (1995) attempts to relate issues such as intonation, focus and the various readings obtained for a given generic sentence.

[16] C.f. *generalised quantifiers* (Barwise and Cooper 1981).

$$GEN[x, s;](x = mary \wedge smoking(x, s) \; ; \; after.dinner(s))$$
$$GEN[x, s;](x = mary \wedge after.dinner(s) \wedge in(x, s) \; ; \; smoking(x, s))$$

where s is a situation, $after.dinner(s)$ denotes that s occurs after dinner, and $smoking(x, s)$ denotes that x is smoking in situation s. TGB contains numerous other examples illustrating the use of GEN to express the meaning of generic sentences. Of course, it should be remembered that GEN does not provide a semantics for generics, since there first needs to be some semantic interpretation of the operator itself. However, it does provide a convenient way of representing the meaning of a generic sentence, particularly when there is more than one possible reading for the sentence.

4.1.4 Default Logical Semantics of Generics

While the GEN operator provides a convenient notation for expressing readings of generic sentences, it does not address the issue of a formal semantics for generics. Several notable attempts at a semantics for generics have been made in the literature, but the various issues described above have not been satisfactorily addressed by any one approach.[17] My main focus in this brief discussion is on semantic accounts of generics and default rules that lead directly to logics for reasoning about generics and with defaults.

4.1.4.1 Traditional AI Approaches to Default Reasoning

There are several formal approaches to default reasoning in the AI literature, the major ones being *Default Logic* (Reiter 1980), *circumscription* (McCarthy 1980; McCarthy 1986) and *autoepistemic logic* (Moore 1985). These approaches are quite different in their slant on the problem of default reasoning, even though there are several results that demonstrate a close correspondence between the three (see (Lukasziwicz 1990) for a survey of such results). Default logic represents a default rule as a rule of inference involving a consistency test. For example, the rule "Birds fly" is represented as follows:

$$\frac{bird(x) \; : \; flies(x)}{flies(x)}.$$[18]

[17] TGB contains a brief overview of several approaches to the semantics of generics, and the associated problems.

[18] Default rules that contain free variables are actually slightly problematic in default logic, but I will ignore this fact for this presentation.

This rule roughly translates as follows: If x is known to be a bird and it is consistent to assume that x flies, then it can be inferred that x flies. Reiter's account of default logic is a purely syntactic one, although there have been some attempts at constructing a model-theoretic semantics for it.[19]

Autoepistemic logic is a modal theory of nonmonotonic reasoning that is basically meant to model an agent's reasoning about its own beliefs. An *autoepistemic theory* is a set of sentences formed from the usual connectives plus a modal operator L, where $L\phi$ is to be interpreted as "ϕ is currently believed". A default rule such as "Birds fly" can then be represented as follows:

$$\forall x(L\, bird(x) \wedge \neg L \neg flies(x) \supset flies(x)),$$

which translates roughly as: If x is believed to be a bird and it is not currently believed that x cannot fly, then x can fly.[20] The semantics of an autoepistemic theory T is given by the notion of a *stable set*—T is *stable* if it satisfies the following:

1. T is closed under logical consequence;
2. if $\phi \in T$ then $L\phi \in T$;
3. if $\phi \notin T$ then $\neg L\phi \in T$.[21]

The conclusions supported by a set of premises A are those that are contained in all stable autoepistemic theories that contain A.[22]

Circumscription is based on *minimising* the extensions of predicates. For simplicity, I will adopt the following notation: given a set of sentences Φ some of whose predicate symbols are in P, $CIRC(\Phi; P)$ is a second-order sentence that minimises the extensions of the predicates in P and allows all other predicates in Φ to vary.[23] The predicates

---

[19] Notably, Rounds and Zhang (1995) use Scott's *domain theory* to provide a semantics whereby a default rule is interpreted as a semantic constraint on allowable *information systems* (which are the models of domain theory).

[20] The reason why all the premises and none of the conclusions must be preceded by the modal operator is not straightforward—see (Konolige 1988) for some explanation.

[21] Moore (1984) also provides a possible-worlds semantics based on $S5$.

[22] See (Moore 1985) for precise details, or (Konolige 1988) for an alternative characterisation.

[23] Circumscription has undergone several revisions since McCarthy's original formalisation. The original presentation in (McCarthy 1980) involved an infinite set of first-order sentences. However, this was reformulated as a second-order sentence in (McCarthy 1986) since the original formulation was incomplete with respect to the minimal-model semantics. A further problem required the introduction of the notion of *variable* circumscription, whereby some predicates that were not minimised could be allowed to "vary" in their interpretation. Following Morreau (1992a), I simply allow all non-minimised predicates to vary, which simplifies the presentation and is sufficient for representing generics. See (McCarthy 1986) or (Etherington 1988) for further details of variable circumscription.

in P are said to be *circumscribed.* The consequences of $CIRC(\Phi; P)$ are perhaps more easily illustrated by the following definition involving *preferred models.*[24]

Definition *Let T be a first-order theory, some of whose predicates are in P, and let M, N be first-order models of T. M is* preferred to *N wrt P, written $M \leq_P N$, if*

1. *M and N have the same domain and assignment to terms;*
2. *if p is an n-ary predicate in P and M satisfies $p(a_1, ..., a_n)$, then so too does N.*

*M is said to be $\leq_P$-*minimal *if there is no model M' such that $M' \leq_P M$ and $M' \neq M$. The consequences of $CIRC(\Phi; P)$ are the sentences that are true in every $\leq_P$-minimal model of Φ.*

Using circumscription, a generic such as "Birds fly" is represented by the theory $CIRC(\Phi; \{ab\})$, where Φ is the following sentence:

$$\forall x(bird(x) \wedge \neg ab(x) \supset flies(x)).$$

This can be interpreted as: "Any bird that is not abnormal flies". Since the predicate ab is circumscribed, the only things considered to be abnormal are those that are logically entailed to be so, and so $CIRC(\Phi; \{ab\})$ supports the conclusion $flies(a)$ for any a such that $bird(a)$ is a consequence of Φ and $ab(a)$ is not a consequence of Φ.[25]

The formal systems briefly described above were primarily designed for defeasible reasoning—i.e. drawing default conclusions about individuals. However, since they are all designed to handle exceptions to default rules, they can be seen as potentially providing a semantics for generics. Default logic has several immediate problems in this respect, emanating from the fact that it represents generics as inference rules. For example, this means that nested defaults cannot be represented, and there can be no logic for reasoning about the generics themselves. However, even the model-theoretic nonmonotonic logics have their problems.[26] For example, the circumscriptive representation of the following generic suffers from the extensionality problem:

[24] Actually, the equivalence of the syntactic version of circumscription with the preferred-models version relies on a (simple) finiteness condition. For details, see (Etherington 1988).

[25] Morreau (1992a) points out a number of problematic issues with the circumscription representation of generics. In particular, when differing ways of being "abnormal" need to be specified, different abnormality predicates effectively need to be used, which puts an onus on the user of the logic that should be taken care of by the semantic framework itself. There are also problems with the way multiple abnormality predicates interact—see (Morreau 1992a) for details.

[26] Since circumscription is the more powerful and more widely proffered as a possible semantics of generics, I will discuss the problems with respect to it. Morreau (1992a) discusses these issues in some detail.

"Elderly members drink for free."

This is represented by the following sentence Φ:

$$\forall x(elderly(x) \wedge member(x) \wedge \neg ab(x) \supset drinks.free(x)).$$

However, suppose there are no elderly members—let ψ be the sentence $\neg\exists x(elderly(x) \wedge member(x))$. It turns out that $CIRC(\psi; \{ab\}) \models \Phi$—i.e. the above generic is true simply by the fact that there are no elderly members at all. A further problem involves what Morreau calls the *Irrelevance* problem. From the premises

$$\forall x(bird(x) \wedge \neg ab(x) \supset flies(x))$$
$$\exists x \neg flies(x)$$
$$bird(tweety)$$

the conclusion $flies(tweety)$ is not obtained, even though the second sentence (i.e., there is some non-flying individual in the universe) seems perfectly reasonable. This is because there are (minimal) models for which the set of individuals is a singleton set, and for such models *tweety* must be the non-flying individual.[27] These specific problems are not so much to do with circumscription. Rather, they concern the use of universal quantification and material implication in an attempt to model what in effect is a generalised quantifier, i.e. GEN. As with other generalised quantifiers, it seems that GEN cannot be adequately defined in terms of standard quantification and implication.[28]

There are also problems with the patterns of reasoning supported by these logics. In particular, none of them supports Specificity. There have been extensions of the basic systems purposely designed to handle Specificity—e.g. *ordered default logic* (Touretzky 1984), *prioritised circumscription* (Lifschitz 1985), *hierarchical autoepistemic logic* (Konolige 1989). However, this sort of extension is exactly what Morreau (1992a) disagrees with so strongly as *ghosts in the machine.* Specificity is, however, a standard pattern of inference supported by default inheritance networks. An inheritance network is a directed acyclic graph, whose nodes are properties and for which an arc represents default inheritance. Although the expressive power of an inheritance network is fairly inflexible, the patterns of inference licensed by the path-based reasoning techniques are

[27] Etherington et al. (1991) combat this problem by basically assigning a context, or scope of individuals, to default rules, which is very much in the spirit of the channel theoretic analysis. Etherington et al. view the notion of scope as fundamental to default reasoning, and use their formal system to tackle several other problematic issues arising in default reasoning.

[28] C.f. Barwise and Cooper's (1981) work on generalised quantifiers.

very powerful.[29] However, default inheritance networks lack a comprehensive semantic interpretation.[30]

4.1.4.2 Logics of Normality

An important recent trend in the semantics of generics has involved what could be called possible-worlds semantics of *normality.* The *normative* view of generics involves the claim that a generic sentence involves universal quantification over the "normal" individuals satisfying the antecedent. For example, the generic "Birds fly" is seen as asserting "All *normal* birds fly".[31] Carlson (1977), for one, has raised several issues that he believes to be problematic for normative treatments of generics. The most problematic issue concerns the question of what constitutes a "normal" individual. For example, it could be that no bird is "normal" in *all* respects. This problem manifests itself in examples such as the following:

> "Peacocks lay eggs."
> "Peacocks have bright feathers."

It is only *male* peacocks that have bright feathers and only *female* peacocks[32] that lay eggs. Morreau (1992a), who is a recent proponent of the normative view, suggests that this problem is more a pragmatic one—the correct logical forms of the above generics should reflect the fact that the quantification in the first is restricted to male peacocks and in the second to female peacocks.[33] However, it does seem acceptable to assert the following:

---

[29] There are some important limitations, however. For example, arcs can only be effectively composed transitively, which rules out Addition of generics (as described above). Also, no conclusions can be obtained using the generics in a contraposed manner (although work by Padgham (1989) is an exception)—while Contraposition is not valid for generics, there are many instances where its use is reasonable. As with other AI formalisms for reasoning with defaults, inheritance networks were primarily designed for drawing default conclusions about individuals, rather than reasoning about the default rules themselves.

[30] There have been some attempts however—e.g. (Bacchus 1989; Boutilier 1989; Stein 1992). More relevantly, Vogel (1995) provides a channel theoretic interpretation of default inheritance networks.

[31] This is similar in viewpoint to the representation of generics using circumscription. However, the following systems are more highly intensional and do not suffer from the particular defects described above in the context of circumscription.

[32] One may want to argue that female peacocks should really be referred to as "peahens"—i.e. the term "peacock" only really applies to the male of the species. However, it is common enough to use the male term for a species to refer to both male and female individuals of that species. In particular, the above sentences certainly seem acceptable.

[33] The same sort of counter is made to combat the problems raised by generics such as:

> "Dutch are good sailors."

where it is only Dutch sailors that are good sailors (i.e. not Dutch farmers).

"Peacocks lay eggs and have bright feathers."

This suggests that Morreau's approach is not such a reasonable solution.

Other potential problems for modal logics of normality, and the normative analysis in general, are discussed by *TGB* and Carlson (1977). Carlson is concerned that the normative analysis may not lead to a "graduated" view of normality—i.e. he shows that normality seems to come in degrees (an individual may be normal in some respects while abnormal in others). In fact, Morreau's logic of generics has problems with this pattern of reasoning. *TGB* discusses examples that seem to require normality orderings that are unusual at best—for example, the sentence

"Turtles are long-lived."

requires a normality ordering whereby the "usual" situation (i.e. a turtle is killed by a predator very early in life) is less normal than one where a turtle avoids its early difficulties and lives on to a ripe old age.

The logics of normality I describe here are obtained from conditional logics. Rather than the worlds-accessibility function being interpreted as a "closest-worlds" function, however, it is seen as accessing "most normal" worlds.[34] While Boutilier (1992) uses an accessibility relation, Morreau employs a worlds-selection function. In each case, the intention is to model an agent's "expectations" about what normally holds—the accessibility relation represents an ordering of normality, while the selection function selects worlds that are "normal" with respect to a given proposition. An important aspect of these approaches is that in neither case is a "global" ordering of normality involved—normality, whether seen as an ordering on worlds or as selecting a set of worlds, is only defined relative to a world.[35]

Boutilier's Logic of Normality

Delgrande's (1988) initial work on defining a logic of default reasoning from a Stalnaker-Lewis logic of conditionals was extended by Boutilier (1992) in his doctoral thesis.[36]

[34] Strictly speaking, one should not talk about worlds that are "most normal" any more than one should talk about worlds that are "closest" (c.f. Lewis' (1973) arguments against Stalnaker's *limit assumption*). However, it does simplify the informal description to talk this way.

[35] This is exactly analogous to "nearness", as used in possible-worlds interpretations of conditional sentences.

[36] Not only does Boutilier's logic allow nested conditionals (and therefore nested default rules) where Delgrande's doesn't, Boutilier also shows that the "flat" (i.e. non-nested) component of his logic contains

Boutilier's basic logic of normality is a conditional logic based on the modal logic *S4*, with the worlds-accessibility relation interpreted as a normality ordering—a model of his logic is a model of (propositional) *S4* with an extension to handle the conditional connective (interpreted as the generic operator), which I will denote by the symbol '$\rightarrow$'.[37] The relevant part of the definition is the following.

Definition *A* CT4-model[38] *is a triple* $M = \langle W, R, \|.\| \rangle$, *where* W *is a set of worlds,* R *is a reflexive, transitive binary relation on* $W \times W$ *and* $\|.\|$ *maps sentences of the logic to* 2^W. *Support for the conditional* $A \rightarrow B$ *is defined as follows:*

$M, w \models A \rightarrow B$ *iff for all worlds* w_1 *such that* wRw_1,

(a) $\exists w_2 : w_1 R w_2,\ M, w_2 \models A$ *and* $\forall w_3 : w_2 R w_3\ M, w_3 \models A \supset B$; or

(b) $\forall w_2 : w_1 R w_2,\ M, w_2 \not\models A$.

As I stated earlier, the worlds-accessibility relation is interpreted as a normality relation—if wRv, then v is at least "as normal" as w. $A \rightarrow B$ is true at a world w if, in the limit of worlds increasingly "more normal" than w, B holds whenever A holds.[39] As a simple example, consider a model M, world w and sentence $bird \rightarrow flies$, and suppose there is a most normal *bird*-world w' (i.e. wRw' in M, $M, w' \models bird$, and there is no world w'' such that $w'Rw''$ and $M, w'' \models bird$)—then $M, w \models bird \rightarrow flies$ iff $M, w' \models flies$.[40]

Boutilier extends *CT4* by adding the extra condition that the accessibility relation be *connected*—i.e. if wRv and wRu then either uRv or vRu. The resulting logic (*CT4D*) is equivalent to the modal logic *S4.3*. *CT4D* is clearly more powerful than *CT4*—in particular, it satisfies the conditional axiom *CV* (a weak form of Monotonicity) and Kraus et al.'s (1990) rule of Rational Monotonicity. In fact, *CT4D* proves to be a very powerful logic of generics. By treating the *assertions* of Kraus et al. as conditional sentences,[41] Boutilier shows that *CT4D* is equivalent to their notional of *rational entailment*, as

all theorems of Delgrande's logic, as well as some that are not supported by Delgrande (Boutilier 1992, Theorem 4.12). I will not describe Delgrande's logic here.

[37] A sentence $A \rightarrow B$ can actually be expressed as $\Box(\Box\neg A \vee \Diamond(A \wedge \Box(A \supset B)))$ in Boutilier's logic, but Boutilier chooses to treat the conditional connective as a primitive.

[38] CT4 is Boutilier's basic conditional logic, based on S4.

[39] A special case of this is when there is some "maximally normal" world w_1 accessible from w at which A holds. In such a case, B must hold at w_1. In the general case, however, there can be an infinite chain of increasingly normal worlds accessible from w.

[40] To be precise, I should also impose the condition that R is connected (to prevent there being other maximally normal *bird*-worlds.)

[41] Kraus et al.'s (1990) results are described in Chapter 5.

well as showing that it is equivalent to Pearl's (1988, 1990) probabilistic ϵ-*entailment*. However, the Irrelevance problem arises when Boutilier attempts to define a system for defeasible reasoning (about individuals) based on his logic of normality—e.g. $flies$ is not a (defeasible) consequence of $bird \rightarrow flies$ and $bird \wedge green$.

Boutilier attempts to overcome the problem of Irrelevance via the concept of *inaccessible* worlds: a world w' is inaccessible from w if it is not the case that wRw'.[42] In particular, Boutilier defines two new modalities reflecting truth in all worlds and truth in all *less* normal worlds, respectively, leading to the logic *CO*. The first modality is used to redefine the conditional connective '$\rightarrow$' in such a way as to account for the property that the truth of a normative statement depends on *all* worlds, not just on normal worlds. The second modality is used to define a new conditional connective '$>$': a sentence $A > B$ means that at most maximally normal A-worlds, *at most* $A \wedge B$ is true. Boutilier proves that for any sentence α that is propositionally consistent with $A \wedge B$, $A \wedge \alpha \rightarrow B$ is a CO-consequence of $(A \rightarrow B) \wedge (A > B)$. This leads to a method for defeasible reasoning that avoids the Irrelevance problem. For example, $(bird \rightarrow flies) \wedge (bird > flies) \models_{CO} (bird \wedge green \rightarrow flies)$. However, Boutilier concedes that the approach has its limitations (Boutilier 1992, pg. 93) and requires a preference relation on models so as to model Lehmann's (1989) system of *rational closure* and Pearl's (1990) system of *1-entailment*—these systems are Lehmann's and Pearl's own extensions to their respective default logics, modified (using extra-logical techniques) in response to the Irrelevance problem.

Morreau's Commonsense Entailment

Morreau's (1992a) thesis is a significant contribution to the normative view of generics. As well as arguing extensively for the normative view as a viable foundation for the semantics of generics, Morreau develops a logic of generics and default reasoning with many attractive properties. One of Morreau's major concerns is for this logic to capture the necessary patterns of reasoning for generics and defaults without any sort of *ad hoc* treatment usually required for handling the more problematic aspects of default reasoning (e.g. ordering default rules so as to invalidate Specificity; solutions to the Irrelevance problem adopted by logics based on conditionals logics). Morreau is adamant

---

[42]In *CT4D*, $w'Rw$ holds for any such pair of worlds since *CT4D* can be characterised by *totally connected* accessibility relations, i.e. relations for which wRw' or $w'Rw$ for all worlds w, w'.

that a logic of generics should be based on independently motivated foundations, and that the necessary properties and patterns of inference required of such a logic should follow from these foundations.[43]

Morreau's logic of *Commonsense Entailment (CE)* is basically a logic of normality incorporating a dynamic revision process. The main function of this dynamic component is *normalisation*—revising an information state s so that it is as "normal" as possible. This process is a very complex one, but is at the heart of commonsense entailment and hence warrants some discussion.

The conditional logic of normality involves assigning a "selection" function to each pair of world and proposition, where a proposition is represented as a set of worlds.

Definition *A* CE-frame *is a triple $F = \langle W, D, {}^{*} \rangle$, where W is a non-empty set of worlds, D is a a non-empty set of individuals, and ${}^{*} : W \times 2^{W} \rightarrow 2^{W}$.*

Basically, ${}^{*}(w, p)$ contains those worlds that support (*wrt* w) all that "normally" holds when p holds. For example, if w is a world that supports the proposition "Birds fly" and p is the proposition that some bird b is a bird, then ${}^{*}(w, p)$ contains only worlds which support the proposition that b can fly. Given that b may not be a "normal" bird, then w may not be contained in ${}^{*}(w, p)$. However, the following condition—i.e. that p normally holds whenever p holds—can be safely imposed on the selection function:

Facticity: ${}^{*}(w, p) \subseteq p$.

Another important condition that Morreau imposes on the worlds-selection function is the following.

Dudley Doorite:[44] ${}^{*}(w, p \cup q) \subseteq {}^{*}(w, p) \cup {}^{*}(w, q)$.

The interaction of this constraint with other properties of the logic leads to a version of Specificity being supported as a valid pattern of inference. Denoting the generic connective by the symbol '$\rightarrow$' (as above), satisfaction of generic sentences is defined as follows.

Definition *A* base CE-model *is a quadruple $M = \langle W, D, {}^{*}, \|.\|_{M} \rangle$, where $\langle W, D, {}^{*} \rangle$ is a CE-frame and $\|.\|_{M}$ interprets predicate symbols. Support for a conditional sentence*

[43] As will be seen below, Morreau is not quite successful—in particular, his logic cannot deal with standard patterns of Specificity.

[44] The strange name for this condition is due to Morreau (1992a). I could not track down his motivation for it.

$\phi \rightarrow \psi$ is defined as follows:

$$M, w \models \phi \rightarrow \psi \textit{ iff } {}^*(w, \|\phi\|_M) \subseteq \|\psi\|_M.$$

This defines a base model in Morreau's system, and a (monotonic) logic of normality is obtained by defining validity and consequence in the standard way. The full system of *commonsense entailment* involves a defeasible notion of consequence, which requires the process of *normalisation.* Morreau describes the process of defeasible reasoning as involving the following abstract steps:

1. Assume the premises in question, and no more;
2. Assume individuals to be "normal", if this can be done consistently;
3. Check if you are then forced to assume the conclusion.

The first step requires updating an information state with a set of sentences (i.e. the premises); the second step involves normalising the resulting state; the final step requires checking whether the normalised information state supports the conclusion.

The basis of an information state is a set of worlds.[45] Updating a state s (relative to a model M) by a formula Φ is via the following *update* operation: $s + \Phi = s \cap \|\Phi\|_M$.

Definition *A* CE information model *(or simply* CE-model*) is a tuple $\langle 2^W, D, ^*, \|.\|, + \rangle$, where $\langle W, D, ^*, \|.\| \rangle$ is a base model. An information model M (based on base model M') and information state s are said to support a sentence Φ, written $M, s \models \Phi$, iff $M', w \models \Phi$ for all $w \in s$.*

Morreau shows that a *canonical* information model can be constructed and the rest of this presentation restricts attention to this model (much the same way that Morreau does), allowing models to be ignored and focus to be restricted to information states.

The process of assuming only a given set of premises Γ involves starting with the empty information state Θ (which can be taken to be the set of all possible worlds—this state satisfies all and only the logical truths) and updating it, via the + operator, with each sentence in Γ. The process of normalising the resulting state s with respect to a proposition p is defined as follows, where ${}^*(s, p) = \bigcup_{w \in s} {}^*(w, p)$:

$$\mathcal{N}(s, p) = \begin{cases} s \setminus (p \setminus {}^*(s, p)) & \text{if } s \cap {}^*(s, p) \neq \emptyset \\ s & \text{otherwise.} \end{cases}$$

[45] This is a fairly standard technique, and is also employed by Veltman (1993).

This operation requires some intuitive explanation. The worlds in $^*(s,p)$ are worlds which are "normal" with respect to p and s—i.e. everything that normally holds when p holds (wrt worlds in s) holds in this information state. If this state, which is somehow "maximally normal" with respect to s and p, is *in*consistent with the current state s, then the normalisation of s is s itself. However, if it *is* consistent, then the normalisation of s is formed as follows: $(p \setminus {}^*(s,p))$ is the set of worlds in which p holds but which are *not* normal wrt s and p; $\mathcal{N}(s,p)$ is formed by removing such worlds from s—i.e. $\mathcal{N}(s,p)$ is s with all the non-normal p-worlds removed. As Morreau puts it, "Normalisation strengthens an information state by adding the assumption that to the extent that some proposition p holds, p holds along with everything which is normally the case where p holds" (Morreau 1992a, pg. 104).

The normalisation process is more complicated than this, however. Since an individual may be a representative of several kinds, each of which is the premise of some generic, the normalisation process must be *iterated* until a fixpoint is reached—i.e. a state s for which $\mathcal{N}(s,p) = s$ for each proposition p. Furthermore, the *order* in which propositions are used in the normalisation process matters, so the normalisation operation has to be iterated in all possible ways—i.e. using all permutations of the pertinent propositions.[46] This leads to the following definition.

Definition *An information state t is* accessible from *a state s just in case there is some ordinal γ and sequence $\{s_\alpha\}_{\alpha \leq \gamma}$ of information states such that:*

1. *$s_0 = s$;*
2. *for each successor ordinal $\alpha < \gamma$, there is a p such that $s_{\alpha+1} = \mathcal{N}(s_\alpha, p)$;*
3. *for limit ordinals $\lambda \leq \gamma$, $s_\lambda = \bigcup_{\alpha<\lambda} s_\alpha$; and*
4. *$s_\gamma = t$.*

Given the discussion above, the state t is of interest when it is a fixpoint of the normalisation operation (with respect to *all* propositions p). The third step in Morreau's outline of default reasoning—i.e. checking whether one is forced to accept the conclusion, having assumed the premises (and no more) and as much "normality" as is consistent—is performed by checking *all* normalisation fixpoints accessible from the state s that consists of the empty state updated by the given set of premises (i.e. $s = \Theta + \Gamma$).

[46] The pertinent propositions are basically the premises (instantiated in their arguments by all possible individuals) of any conditional sentence in a relevant theory Γ of commonsense entailment—the following definitions should be relativised to such a Γ, but I have glossed over this for the sake of simplicity.

Definition *A set of sentences* Γ *is said to* commonsense entail ϕ, *written* $\Gamma \approx\!\!\!\!| \ \phi$, *iff for* any *fixpoint* s *accessible from* $\Theta + \Gamma$, $\mathcal{M}, s \models \phi$ *(where* $\mathcal{M}$ *is the canonical model).*

The process of normalisation is a powerful one and overcomes the problems of Irrelevance so prevalent in other approaches to default reasoning based on conditional logics of normality. Normalisation effectively involves performing a consistency check on an information state s and then assuming s to be "as normal as possible". Any information that is "irrelevant" to some premise p cannot prevent the inference of a conclusion that is deemed to normally follow from p. In particular, we have the following entailment:

$$\forall x(bird(x) \rightarrow flies(x)) \mid\!\approx \forall x(red(x) \wedge bird(x) \rightarrow flies(x)).$$

Appendix C.1 shows in detail how this entailment is obtained in commonsense entailment, thereby illustrating the process described above.

Morreau's logic of generics is a powerful one and leads directly to a logic of default reasoning (i.e. allowing defeasible conclusions about individuals to be drawn). The possible-worlds framework ensures that it is appropriately intensional, avoiding the problem suffered by circumscription whereby asserting the fact that there is some non-flying individual prevents the inference that the bird *tweety* flies (as discussed above). Morreau (1992a) also shows that many of the desirable patterns of reasoning, such as defeasible *modus ponens* (including the case where one generic is nested inside another) and a weakened form of Transitivity, are supported. When the Dudley Doorite constraint is imposed on the normal-worlds selection-function, the resulting logic satisfies the following weakened form of Specificity, called *Taxonomic Specificity:*

$$\begin{array}{l} \phi(t) \\ \forall x(\phi(x) \supset \psi(x)) \\ \forall x(\psi(x) \rightarrow \tau(x)) \\ \forall x(\phi(x) \rightarrow \neg\tau(x) \\ \hline \neg\tau(t) \end{array}$$

Of course, this is isomorphic to the standard Penguin example, and (correctly) allows the inference that a given penguin cannot fly. This is an important property. However, if the material implication in the above is replaced by the generic conditional (i.e. $\forall x(\psi(x) \supset \psi(x))$ is replaced by $\forall x(\phi(x) \rightarrow \psi(x)))$, thereby obtaining standard Specificity, then Morreau's logic is no longer able to make the inference that $\neg\tau(t)$ holds.

However, as shown by the *student* example (see Section 4.3.5), this pattern of reasoning also seems intuitively valid. Morreau (1992a) discusses a scheme whereby normalisation can be *prioritised,* but this is simply introducing a "ghost in the machine"—exactly analogous to that which he so strongly criticises in circumscription and other systems that require meta-logical additions in order to capture intuitively valid patterns of reasoning. Morreau's basic logic also does not support Graded Normality—in fact, one of Carlson's (1977) original qualms about the normative interpretation of generics was that it would have problems with the concept of normality coming "in degrees".

Morreau's logic of Commonsense Entailment is clearly powerful and captures many intuitive properties of generics and their valid patterns of reasoning, as well as providing a fairly comprehensive account of generics within the viewpoint that generics quantify over normal individuals. However, the model-theory associated with Commonsense Entailment is extremely complex, and the fact that such central patterns of inference as Specificity and Graded Normality are not validated seems to suggest that the account is lacking in some crucial way.

4.2 A Channel Theoretic Analysis of Generics

Generic sentences describe regularities: this claim is the basis for the channel theoretic model of generics. Of itself, this claim is uncontroversial enough—for example, *TGB* describes a generic sentence as stating a "law-like regularity". The main difficulty in defining a semantics for generics has generally involved the attempt to provide a definition of regularity that caters for exceptions and the other issues raised in Section 4.1.2. This issue—of reducing the concept of regularity to other "primitive" concepts—is, of course, not addressed by channel theory: in channel theory, a regularity is a channel-type and is taken to be a primitive concept of the theory, as discussed earlier. However, channel theory does provide features in its account of the nature of regularity which are important in relation to the problematic issues in the semantics of generics.

4.2.1 A Simple Semantics for Generics

A generic describes a regularity; a channel supports a regularity. The channel theoretic analysis of conditionals of the previous chapter involved taking a channel to be the demonstrative content of a conditional. The channel theoretic analysis of generics involves a similar move—the demonstrative content of a generic sentence is taken to be a channel while the descriptive content is taken to be a channel-type, or constraint (i.e. a regularity). The demonstrative content associated with a generic seems to play an important role in some of the problematic issues in the analysis of generics, as is discussed below. The intensional nature of channel-types means that other problematic issues—to do with the independence of the truth of a generic from the properties of associated individuals—are avoided. In particular, this leads to a view of generics being *simply true* (in a situated way) in the manner described in Section 4.1.2—i.e. of themselves, rather than because of the properties of the associated individuals. This requires an assumption of *maximal normality* before a logic that uses generics for defeasibly reasoning about the properties of individuals can be obtained—this is the topic of Chapter 5.

The most important property of channels that makes them suitable as a basis for a semantic analysis of generics is the fact that they allow exceptions. In the channel theoretic model, an exception (using the term informally) to a regularity can occur in one of three ways: by not being part of the classifications with which the channel $\mathcal{C}$ is concerned; by corresponding to a token that is a pseudo-signal in $\mathcal{C}$; by corresponding to a connection that is an *exception*[47] in $\mathcal{C}$. These different possibilities, and the various sorts of exceptions they seem to account for, are discussed in Section 4.2.2, as is the channel theoretic model's behaviour with respect to the other issues raised in Section 4.1.2.

The channel theoretic analysis of generic sentences is very similar to the analysis of conditional sentences presented in Section 3.1.1. In particular, a generic sentence involves a claim about the structure of a particular channel. In the case of a generic, however, the claim only concerns the type-level regularity supported by the channel.

Definition *The content of a generic sentence is an Austinian proposition of the form* $(\mathcal{C} : \gamma)$*, where* $\mathcal{C}$ *is a channel and* γ *is a constraint (i.e. a channel-type). If* $\mathcal{G}$ *is a classification that classifies channels by constraints, then for* $\mathcal{C} \in tok(\mathcal{G})$ *and* $\gamma \in typ(\mathcal{G})$*,* $(\mathcal{C} : \gamma)$ *holds in* $\mathcal{G}$ *iff* $\gamma \in typ(\mathcal{C})$*.*

As can be seen from the definition, whether or not a channel $\mathcal{C}$ supports a generic is

[47] It is unfortunate that the term "exception" is used in channel theory in this technical sense. When I talk of an exception to a regularity, I will typically be using the term in its informal sense—i.e. meaning some sort of counter-example to a general rule. If I want to get at the technical meaning of the term, I will usually explicitly refer to a *connection* being an exception. Hopefully, this will resolve any potential confusion.

independent of the properties of individuals related to that generic—all that matters is whether or not $\mathcal{C}$ contains the appropriate type (i.e. supports the associated regularity). While this skirts the most problematic issue in the analysis of generics (or regularities for that matter)—i.e. providing a reductionist account for them—one can take the viewpoint that the channel theoretic model is consistent with any particular account of regularity one may wish to provide. However, the standard situation theoretic and channel theoretic view of regularities is that they are irreducible components of the world and should therefore be primitive objects of the theory.

As a simple illustration of the above idea, consider the archetypal example of a generic: "Birds fly". This sentence makes a claim about some[48] channel $\mathcal{C} : B \Rightarrow F$ that $typ(\mathcal{C})$ contains a constraint of the form $bird \rightarrow flies$. Given a classification $\mathcal{G}$ which classifies channels by channel-types (i.e. constraints), $(\mathcal{C} : bird \rightarrow flies)$ will hold in $\mathcal{G}$. Note that $\mathcal{C}$ supports $bird \rightarrow flies$ regardless of the properties of individual bird-tokens—in particular, some bird-tokens will not be contained in $tok(B)$, some will be pseudo-signals in $\mathcal{G}$, some will be involved in connections (of the form $b \mapsto b$) that are exceptions in $\mathcal{C}$, while birds that are "normal" (with respect to the above generic) will be involved in connections in $\mathcal{G}$ that are classified by the constraint $bird \rightarrow flies$.

As mentioned earlier, this analysis makes generics *simply true* or *simply false.* The question that arises is: *what* is it that makes a generic sentence true/false? *B&S* take the view that (certain) regularities are part of the natural order of things. Under this view, a generic is true with respect to a channel $\mathcal{C}$ if it describes a regularity that actually holds in the world (at least, within the context of $\mathcal{C}$). This is actually not of great importance to the current enterprise. Since this chapter is primarily about reasoning with generics, then my primary concern is with consequence—i.e. given the assertion that some channel $\mathcal{C}$ supports some regularity γ, what other information can be derived? In particular, what other generics can be inferred given the generic that asserts that $(\mathcal{C} : \gamma)$ holds? I turn to the topic of a logic of generics in Section 4.3.

An account of generics related to the one described here is provided by ter Meulen (1986). Ter Meulen discusses the use of situation theoretic constraints as the basis of a semantics for generics, focussing on the way in which what she calls *recalcitrant* situations—i.e. situations that are counter-examples to a generic—do not necessarily

[48]Exactly *which* channel the sentence makes a claim about is, of course, not determined by the channel theoretic model. This issue was discussed (in the context of conditional sentences) in Section 3.1.1.

undermine the general rule (represented by a constraint). As here, a constraint is taken to be a primitive entity, with its truth being independent of the properties of the associated individuals. Similarly to Barwise's (1986) account of conditionals, ter Meulen's analysis of generics can be seen as a precursor to the channel theoretic analysis, using the notion of conditional constraint. The semantics of a generic sentence has no demonstrative content in ter Meulen's account, of course, and ter Meulen also does not provide any sort of logic of generics—her main concern is in providing a situation theoretic analysis of generics that accounts for exceptions.

4.2.2 Properties of the Channel Theoretic Semantics

In this section, I discuss some of the important properties of the channel theoretic analysis of generics, focussing particularly on the issues raised in Section 4.1.2.

Exceptions

An exception to a generic is an individual that satisfies the antecedent of the generic but does not satisfy the consequent. This may occur in one of three ways in the channel theoretic model. It should be stressed that the choice as to which of the three categories an exception falls into is a subjective one—an utterer of a generic sentence may make her utterance based on a model that represents an exception in one way, while the hearer of the utterance may construct a model that represents the "same" exception in a different way.[49] The choices involved in the representation of exceptions are further discussed in Chapter 5.

Token is not contained in the antecedent classification.
Let $\mathcal{C} : C \Rightarrow L$ be the channel about which the generic "Chickens lay eggs" makes a claim. This sentence asserts that $\mathcal{C}$ contains a regularity γ whose antecedent classifies tokens as *chickens* and whose consequent classifies tokens as *egg.layers*. There are many exceptions to this generic of course, such as sick chickens, baby chickens and very old chickens. However, there is a noticeable large class of chickens that constitute counter-examples, namely, male chickens. It could be argued that the class of male chickens is not meant to be covered by the above generic—i.e. only female chickens

[49] A formal account of this form of relativism could be constructed using the channel theoretic model of dialogue of Healey et al. (1993, 1994).

are relevant. This viewpoint is provided by the channel theoretic model, whereby the classification C contains only *female* chicken tokens—i.e. for each $t \in tok(C)$, t is a female chicken. Of course, there is no commitment to such a view—there may be other channels $\mathcal{C}' : C' \Rightarrow L'$[50] involving the same regularity but for which $tok(C)$ contains some male chickens—such chicken tokens are either pseudo-signals or are contained in a connection that is an exception in $\mathcal{C}'$.

Declerk (1991) raises a similar point when he discusses the way in which generics quantify over "relevant" individuals. Note, however, that restriction to relevant individuals cannot be a *sufficient* condition for supporting a generic claim—"Birds don't fly" does not seem supported if attention is restricted to penguins and emus only. The channel theoretic analysis still requires a regularity to hold between the relevant types for the corresponding generic sentence to be supported.

Token is a pseudo-signal in the channel, *or*

Token is involved in a connection that is an exception in the channel.

The question relating to whether a counter-example to a generic should be treated as a pseudo-signal or an exception in a channel is not an easy one to answer. In attempting to address it, it seems as if epistemic issues must be addressed—i.e. if an individual is "expected" (by the agent being modelled) to contradict the general rule, then that individual should correspond to a pseudo-signal; otherwise, the individual should correspond to an exception in the channel. For example, consider again the generic "Chickens lay eggs" and the channel $\mathcal{C}' : C' \Rightarrow L$ from above, where C' contains chickens that do not lay eggs (as well as the usual chickens that do). A male chicken c, where $c \in tok(C')$, seems most naturally represented as a pseudo-signal in this channel. However, a female chicken c', where $c' \in tok(C')$, that is sick and therefore temporarily a non-egg-laying specimen seems more naturally represented as an exception—i.e. there is a connection $c' \mapsto c' \in tok(\mathcal{C}')$, but $c' \mapsto c'$ does not fall within the domain of the regularity described by the given generic.

I stressed in Section 2.3.4 that whether a particular counter-example is best viewed as an exception or as a pseudo-signal is very much dependent on the specific view that is being modelled. In the channel theoretic model of generics, the difference does not in any way matter—the sole concern is that exceptional individuals (i.e. counter-examples to a generic) can be naturally represented within the channel theoretic framework. In a more

[50] In general, $\mathcal{C}$ will be a subchannel of $\mathcal{C}'$.

comprehensive treatment of generics in a dialogue situation, particularly if dealing with speaker/hearer differences and the way each dialogue participant handles a counter-example, the distinction between exceptions and pseudo-signals may take on greater importance. This is particularly so because of the way default inference is defined in Chapter 5—an individual that is "known" to be a counter-example to a default rule (e.g. a bird that is known to be a penguin) must be represented as a pseudo-signal to prevent the rule from being applied to that individual. The task of developing a model of dialogue involving generic sentences is left as a topic for future research.

Generics and Truth

In Section 4.1.2, I briefly discussed Morreau's (1992a) arguments for taking the content of a generic to be a proposition rather than a rule of inference. In the channel theoretic model, the content of a generic is taken to be an Austinian proposition, i.e. the proposition that some given channel supports a specified regularity. The fact that this proposition is endowed with a demonstrative content does not detract from the advantages enjoyed by propositional accounts of generics over inference-rule accounts. In particular, nested generics are handled by this analysis in much the same way as nested conditionals were handled in Section 3.1.2—i.e. by allowing the classification linked by a channel to be the classification $\mathcal{G}$ from the definition in Section 4.2.1.

The way in which the channel theoretic model handles nested generics can be illustrated by once again considering the following example:

> "Current medical theories predict that smoking leads to cancer."

It is not the case that every medical theory makes the above prediction—for example, some medical theories are not concerned with such issues (e.g. those concerned with the treatment of malaria), while research funded by a tobacco company may lead to a theory that smoking is in no way linked to an onset of cancer. Similarly, smoking does not always lead to cancer. This nested generic can be represented in the channel theoretic model by way of a channel $\mathcal{C} : M \Rightarrow \mathcal{G}$, where M classifies theories as being medical theories. The descriptive content of the generic is a regularity $\gamma \in typ(\mathcal{C})$ whose antecedent is a type $medical.theory$ that classifies medical theories and whose consequent is another regularity $\gamma' \in typ(\mathcal{G})$, where γ' is a regularity linking a type that classifies individuals as being smokers to a type that classifies individuals as having cancer (or as being future cancer sufferers).[51]

Pseudo-signals or exceptions in the channel $\mathcal{C}$ are medical theories that do not predict the link between smoking and cancer. For example, if $(t : medical.theory)$ holds in M but there is no connection in $tok(\mathcal{C})$ involving t, then t is a pseudo-signal in $\mathcal{C}$—i.e. t is not related to a channel that supports a smoking–cancer regularity (e.g. if t is not concerned with cancer or smoking-related diseases). Alternatively, if $t \mapsto \mathcal{C}'$ is an exception in $\mathcal{C}$, then t is linked (in $\mathcal{C}$) to a channel $\mathcal{C}'$, but $\mathcal{C}'$ does not support the smoking–cancer regularity. Now suppose $t \mapsto \mathcal{C}'$ is classified in $\mathcal{C}$. By the Principle of Harmony (Section 2.3.3), $\mathcal{C}'$ must support the regularity that smokers are prone to cancer—i.e. $(\mathcal{C}' : \gamma')$ holds in $\mathcal{G}$. However, $\mathcal{C}'$ may itself contain exceptions—even though the general prediction (that smoking leads to cancer) is made, it need not predict that *every* smoker will contract cancer (or indeed that *any* smoker will contract cancer). This example suggests that the channel theoretic analysis allows a coherent treatment of nested generics. However, as was done with conditionals, I will ignore nested generics for the rest of this chapter. The important point is that the channel theoretic framework allows an interpretation of nested generics without requiring any extension.

As mentioned earlier, the channel theoretic analysis leads to a view whereby generics are simply true or simply false. This is because the channel-types (i.e. regularities) are taken to be primitive objects of the theory. The consequences for the channel theoretic model as a system for default reasoning are discussed in Chapter 5. However, the view of regularities as primitive also means that the extensionality problem of Section 4.1.2 is trivially avoided. Consider again the example from before:

> "Elderly members drink for free."

This generic is supported in a channel $\mathcal{C}$ iff $\mathcal{C}$ contains a type γ that corresponds to the appropriate regularity. It makes no difference whether or not there are any elderly members of the club; what matters is whether or not there is a regularity of the required form. The issue that is not addressed is what it is that makes this regularity hold—whether it is some rule of the club, or whether it is "normal" practice of the club's barman not to charge elderly members for their drinks. This is not the concern of channel theory—regularities are part of the fabric of the world and need not (indeed, perhaps can not) be reduced to more primitive concepts.

[51] This is probably over-simplistic, but is sufficient for the purposes of illustration.

The Role of Context

In Section 4.1.2, I briefly discussed how some generics seem to be asserted with respect to a particular background or context.[52] In the channel theoretic analysis, a context is provided by the channel that plays the demonstrative content of a given generic. The role of context alleviates the problems with the examples given in Section 4.1.2. For example, "Chickens lay eggs" can be considered to be restricted to a context in which the only tokens are healthy female chickens of an appropriate age—i.e. the utterer of this generic would intend a channel satisfying this property as the demonstrative content of his utterance. Examples such as "Dutch make good sailors" are handled similarly—in this case, the intended context is one such that only sailor-tokens undergo consideration.

This simple idea seems to offer the potential of quite a powerful analysis, especially when conflicting generics are involved. For example, consider the conflicting pair of generics involving peacocks from Section 4.1.4. These can be assertible at the same time, yet involve disjoint sets of individuals (i.e. only male peacocks have bright tail feathers and only female peacocks lay eggs). A speaker that asserts both would (one assumes) be talking about different channels, each supporting a different regularity about (certain) peacocks. Each channel would be "reliable" in that it need not contain any connections that are exceptions, but each would have a number of pseudo-signals—i.e. the *bright.feathers* channel (call it $\mathcal{C}_f$) would have all the female peacocks as pseudo-signals while the *egg.layers* channel (call it $\mathcal{C}_e$) would have all the male peacocks as exceptions. Now consider the following generic, which I consider to be acceptable:

> "Peacocks lay eggs and have bright feathers."

This seems to be an assertion regarding the channel $(\mathcal{C}_f \parallel \mathcal{C}_e)$—i.e. the parallel composition channel. This channel would have lots more pseudo-signals—in fact, one may argue that all *peacock* tokens are pseudo-signals in this channel since no peacock both lays eggs and has bright feathers. However, the above generic is still supported by this particular channel. Normative approaches to the truth of generics, however, have problems with examples of this sort—since the normative analysis reduces generics to quantification over normal individuals, generics of the above form are universally invalidated since no "normal" peacock both lays eggs and has bright feathers.

[52] "Context" here is taken to mean the channel which the generic describes and not the context the speaker finds herself in when she asserts the generic.

The role of the demonstrative content in the channel theoretic analysis seems to offer some very interesting insights into the information carried by generics and the way they are used. In particular, it seems that an interesting account of disagreement between speaker and hearer as to the channel described by a generic (thereby resulting in misunderstanding in the tokens to which a generic is applicable) could be accommodated in a channel theoretic model of dialogue (involving generics) along the lines of (Healey et al. 1993; Healey and Vogel 1994). This is discussed in Section 6.2 as a topic for future research.

Patterns of Reasoning

The rest of this chapter is dedicated to defining a channel theoretic logic for reasoning with generics. A number of patterns of reasoning, including the test principles of Asher and Morreau (1991), are examined below and found to be satisfied by the channel theoretic logic. Such patterns include Graded Normality and Specificity. Graded Normality is supported since different generics involving the same premise may yet be supported by different channels—an individual that is an exception to the general regularity of one channel need not be an exception to a regularity supported by the other. Specificity, in particular, is considered an important test principle for logics of defeasible reasoning since it proves so troublesome to the traditional model-theoretic accounts from the AI literature.

4.3 A Channel Theoretic Logic of Generics

In this section, I define a system for reasoning about generics, based on the channel theoretic system for reasoning about conditionals of the previous chapter.

4.3.1 From Logics of Conditionals to Logics of Generics

One of the starting premises for the work described in this section is that a system for reasoning with generics and defeasible rules can be based on a system for reasoning about conditionals. In the previous chapter, I developed a system for reasoning about conditionals based on the use of a hierarchy of channels to capture the notion of "implicit background conditions". I take this framework for conditional logic as the starting point for the logic of generics and default reasoning.

Previous work in the AI and philosophical literature investigating the relationship between conditional logic and the logic of generics includes that of Delgrande (1988), Boutilier (1992) and Morreau (1992a). The first two of these suffer from a problem with "irrelevant conditions"—standard logics of conditionals are more sensitive to implicit background assumptions than logics of default reasoning. For example, given the conditional "If tweety is a bird then tweety can fly", one cannot then infer "If tweety is a red bird then tweety can fly" in standard possible-worlds conditional logics. However, given the generic "Birds fly", we certainly expect to be able to infer "Red birds fly" (unless, of course, it is known as a rule that red birds happen to have non-flight qualities). Delgrande and Boutilier each introduces a modification to his logic in order to address this problem.

The channel theoretic system for reasoning about generics or default rules (defined below) avoids the Irrelevance problem by appealing to the following fundamental observation regarding the relationship between conditionals and generics. By restricting just what can be encoded as a background assumption and therefore be used to block certain inferences, the channel theoretic system of conditional logic can be strengthened. In particular, a condition that is "irrelevant" to a given regularity will not be encoded as a background assumption to that regularity.

> **Fundamental Observation:** *Reasoning about generics is basically the same as reasoning about conditionals, under the added proviso that all background assumptions are implicitly contained in the given set of regularities.*

As an example, consider the standard example demonstrating the failure of Monotonicity for conditionals:

> "If there is sugar in the coffee then it tastes good"
> "If there is sugar and diesel-oil in the coffee then it tastes good".

It should be possible to assert the first conditional above without necessarily inferring the second. This turns out to be the case in the conditional logic of Chapter 3 due to the (assumed) implicit background condition that the coffee tastes good if there is sugar in it *and* there is no diesel-oil in it. In a standard default logic, however, the second sentence above is a consequence of the first, in the absence of information that forces a conclusion to the contrary.[53] However, the addition of the sentence

> "If there is diesel-oil in the coffee then it tastes bad"

should be enough to block the undesired default inference.

The above example illustrates the way in which a system for reasoning about generics is obtained from the channel theoretic conditional logic. Given a set of generics, or default rules, a channel hierarchy is *constructed* from them. Explicit and implicit counters to the given generic rules are extracted and incorporated into the channel hierarchy in such a way that the resulting conditional logic supports reasonable patterns of inference for generics. Herein lies the inherent difference between conditional and default reasoning—the latter is simply conditional reasoning under the added assumption that all the information relevant to the hierarchy is explicitly represented in the initial set of defaults rules. Any conclusions that can be drawn, given the represented information, can be feasibly drawn by default. Conditionals, however, invalidate some inferences even if there is no *explicit* background assumption countering the inference. As such, standard logics of conditionals are weaker than is appropriate for a logic of defaults or generics; in particular, this leads to the problem of Irrelevance experienced by Delgrande and Boutilier. The Irrelevance problem is avoided by the channel theoretic system since (by the Fundamental Observation above) irrelevant conditions cannot be used in the construction of an appropriate channel hierarchy for a given collection of generics.

The definition of an appropriate channel hierarchy for a given collection of channels (supporting a corresponding collection of generics) is based on identifying generics whose consequents mutually conflict—e.g. if two constraints are of the form $\phi \to \psi$ and $\tau \to \neg\psi$, then $\neg\tau$ is inserted as a background assumption to the former while $\neg\phi$ is inserted as a background assumption to the latter.[54] This is similar to Pollock's (1987) use of *defeaters* of a default rule, i.e. propositions whose truth undermine a default inference. For example, τ is a *rebutting defeater* to $\phi \to \psi$ above since if τ holds, then it provides evidence for doubting that ψ can be inferred from ϕ. The relation between Pollock's rebutting defeaters and the appropriate channel hierarchy for a system of generics is roughly as follows: if τ is a defeater for $\phi \to \psi$ then $\neg\tau$ is inserted as a background assumption to $\phi \to \psi$. Pollock also considers *undercutting defeaters* to be an important

[53] If this seems counter-intuitive, it is only because we are biased into rejecting the second sentence due to an *a priori* belief that coffee with diesel-oil in it tastes bad.

[54] Actually, the construction is more complicated than this, as is seen below.

(though traditionally neglected) concept in a logic of default reasoning—an undercutting defeater to $\phi \rightarrow \psi$ is a proposition that provides a reason for denying that $\phi \rightarrow \psi$ is a regularity. Undercutting defeaters are ignored in the current setting—in particular, I would need to be able to represent the negation of generics to be able to handle undercutting defeaters.[55] Further investigation of the relation between the constructive definition of a channel hierarchy (below) and Pollock's framework of default reasoning is a possible topic for future research.[56]

4.3.2 Notation and Simplifications

Before presenting the definition of an appropriate channel hierarchy for a given collection of generics, I review some of the simplifications and notational conveniences made earlier, as well as more explicitly define others that are tacitly introduced above.

Firstly, as I have done so far, I will tend to identify connections and constraints with their endpoints (i.e. the tokens and types they connect) and will tend to consider only reflexive channels and signalling relations. Also continuing from the previous chapter, I will tend to ignore the token level of channels—this level plays effectively no part in the system for reasoning about generics (although it assumes greater importance in Chapter 5). In particular, in the definition of a channel hierarchy, I assume existence of superchannels (which encode the background assumptions of given generics) without specifying the signalling relations of such channels. These signalling relations are partially constrained by the subchannel definition, but are otherwise irrelevant to the logic of generics.[57]

Secondly, the representation of generics within the channel theoretic framework is simplified via a couple of assumptions: I assume that each generic is supported by a separate channel and that there is a single classification A which all such channels connect. The second of these assumptions helps to "decontextualise" the logic of generics, which is useful when ensuring that certain patterns of reasoning are validated.[58] Also, I will

[55] Since any channel is itself a classification, this could possibly be done in the extended framework alluded to in Appendix A.2.

[56] Also of interest is Nute's (1993) *Defeasible Prolog,* a computational interpretation of default reasoning based on Pollock's ideas.

[57] The maximal normality condition of Chapter 5 constrains these signalling relations further, in the context of using the associated channels to (defeasibly) reason about the properties of individuals.

[58] For example, if a pair of generics were respectively supported by channels $\mathcal{C} : A \Rightarrow B$ and $\mathcal{C}' : A' \Rightarrow B'$ where either $A \neq A'$ or $B \neq B'$, then these channels could not be composed so as to obtain a channel

treat all regularities supported by channels as defeasible—any "strict" or "taxonomic" informational relationships (i.e. those that do not have any possible counter-examples) are represented via the appropriate type-entailment relation and as a constraint in a logical channel.[59]

Finally, the current framework does not treat complex expressions involving generics, such as conjunctions of generics, negations of generics and nested generics. Of course, since a channel is itself a classification, such expressions could be modelled in the extended classification theoretic framework alluded to in Appendix A.2.

4.3.3 Constructing a Channel Hierarchy

A logic for reasoning about generics is obtained from the logic of conditionals of Chapter 3 simply by defining a method for constructing a channel hierarchy from a given collection of channels, and it is this task to which I now turn. The choice of hierarchy is important—as was discussed in Chapter 3, changing the hierarchy alters the patterns of inference obtained. Rather than exhaustively analysing all possible variations on reasonable patterns of reasoning involving generics, I present a simple definition of channel hierarchy and show that the logic of generics thereby obtained supports the desirable properties and patterns of inference discussed in Section 4.1.2.[60]

Given an initial collection of generics (represented by a collection of channels and the regularities they contain), background assumptions for these generics are obtained simply by considering which rules have conflicting conclusions. Basically, if $\phi \rightarrow \psi$ and $\tau \rightarrow \neg\psi$ are regularities applicable to the same connection, then $\neg\tau$ is a background assumption to the first and $\neg\phi$ is a background assumption to the second. There is a slight qualification needed, however—if, say, individuals that are of type ϕ are generally also of type τ (i.e. there is a regularity of the form $\phi \rightarrow \tau$), then $\neg\tau$ cannot be a background assumption to the first regularity since an individual t being a ϕ carries the information (at least potentially) that it is also a τ. For example, while $\neg$*penguin* can be considered a background assumption behind the regularity $bird \rightarrow flies$, one would not consider $\neg bird$ to be a background assumption behind the

supporting the obvious new generic.

[59] I return to this particular issue in Section 4.4.2.

[60] In this chapter, I only consider patterns of inference involving generics, or default rules. Inference involving defeasible reasoning—i.e. using default rules to draw conclusions about individuals—is discussed in the following chapter.

regularity *penguins*$\rightarrow\neg flies$ since it is nothing to do with a penguin's *bird*-ness that prevents it from flying. In general it is not only the given regularities that have to be checked for conflict but also serial compositions of them—even if $\tau\rightarrow\neg\psi$ is obtained by transitively combining other regularities, then $\neg\tau$ should be inserted as a background assumption to $\phi\rightarrow\psi$. For example, if *king.penguin*$\rightarrow$*penguin* is used to obtain *king.penguin*$\rightarrow\neg flies$, then $\neg$*king.penguin* needs to be inserted as a background assumption to *bird*$\rightarrow flies$. However, in this example, transitive closure also obtains the regularity *king.penguin*$\rightarrow flies$, but $\neg$*king.penguin* should not be inserted as a background assumption to *penguin*$\rightarrow\neg flies$. The fact that a king penguin is a penguin is not a sufficient condition to block $\neg$*king.penguin* being inserted as a background assumption here—for example, penguins are birds but $\neg$*penguin* is still a reasonable background assumption to *bird*$\rightarrow flies$. However, the link between *king.penguin* and $flies$ is obtained *only* via transitive composition with *penguin*$\rightarrow$*bird*—it is for this reason that $\neg$*king.penguin* is ruled out as a reasonable background assumption to *penguin*$\rightarrow\neg flies$.

To achieve full generality in the following definition, the given set of channels should also be closed under applications of parallel composition and contraposition before determining the defeaters. However, this complicates the definition and for the meantime I restrict attention to serial composition, which is enough to (theoretically, at least) model the power of default inheritance networks. I discuss the application of contraposition and parallel composition in Section 4.4.

Definition *Let $\mathcal{S}$ be a collection of channels containing a logical channel $\mathcal{L}$.*[61] *Let $\mathcal{T}(\mathcal{S})$ be the collection of channels obtained by closing $\mathcal{S}$ under standard serial composition. The* kernel default channel hierarchy for $\mathcal{S}$ *is the smallest hierarchy $\mathcal{K}$ of channels containing $\mathcal{S}$ such that:*

> *for each channel $\mathcal{C}_1 : \phi\rightarrow\psi$ in $\mathcal{S}$ and each channel $\mathcal{C}_2 : \tau\rightarrow\psi'$ in $\mathcal{T}(\mathcal{S})$ such that $\psi \perp \psi'$, $\mathcal{K}$ contains a channel $\mathcal{C'}_1 : (\phi \wedge \neg\tau)\rightarrow\psi$, where $\mathcal{C}_1 \sqsubseteq_f \mathcal{C'}_1$ and $f(\phi\rightarrow\psi) = (\phi \wedge \neg\tau)\rightarrow\psi$;* unless *either*
>
> 1. *there is a channel $\mathcal{C} : \phi\rightarrow\tau$ in $\mathcal{T}(\mathcal{S})$;* or
> 2. *$\mathcal{C}_2$ is formed by a composition involving a channel $\mathcal{C}' : \phi\rightarrow\sigma$, for some σ.*

Note that the hierarchy obtained via the above definition is only a "kernel"; a full default hierarchy is obtained by closing such a kernel under the channel operations—

[61] For the rest of the chapter, I will tend to neglect to mention $\mathcal{L}$ and simply take its presence for granted.

this is defined below. The above definition is, I believe, fairly straightforward (given the discussion before) except for the second proviso, which is illustrated by an example below. A default channel hierarchy is meant to capture all the background conditions of a given set of channels *with respect to the information given*—i.e. not *all* feasible background assumptions are captured by this hierarchy, but only those that can be inferred from the initially given collection of regularities. This is the sort of behaviour required of a logic of generics or default rules.

To illustrate the properties of the above definition, consider the standard *penguin* example:

> "Birds fly."
> "Penguins are birds."[62]
> "Penguins don't fly."

These generics are represented by the following collection $\mathcal{S}$ of channels and associated constraints:

> $\mathcal{C}_1 : bird \rightarrow flies$
> $\mathcal{C}_2 : penguin \rightarrow bird$
> $\mathcal{C}_3 : penguin \rightarrow \neg flies.$

By definition, the kernel default hierarchy for $\mathcal{S}$ contains a channel $\mathcal{C'}_1 : (bird \wedge \neg penguin) \rightarrow flies$ such that $\mathcal{C}_1 \sqsubseteq \mathcal{C'}_1$—i.e. $\neg penguin$ is an implicit assumption behind the constraint $bird \rightarrow flies$. Notice that $\neg bird$ is *not* an implicit assumption behind the constraint $penguin \rightarrow \neg flies$—it can't be that non-bird penguins are the ones that can't fly since penguins are birds to start off with! This is captured by the first proviso in the definition—i.e. the kernel default hierarchy does not contain $\neg bird$ as an implicit assumption to $penguin \rightarrow \neg flies$ because $\mathcal{C}_2$ contains the constraint $penguin \rightarrow bird$.

Now suppose we add the generic

> "King penguins are penguins."[63]

to the above collection of generics and the channel and constraint

> $\mathcal{C}_4 : king.peng \rightarrow penguin$

[62]Actually, this generic is taxonomic and does not have any exceptions. Given the discussion in Section 4.3.2, it could be represented as a constraint in the logical channel. However, so as not to detract from the illustration, I simply represent it as a normal regularity—it makes no difference to this example (although it may do for others, as is illustrated in the Multiple Inheritance example in Section 4.3.7).

[63]This is a taxonomic assertion rather than a generic, but again this is ignored for illustrative purposes.

to $\mathcal{S}$. The corresponding kernel hierarchy extends the previous one by the addition of a channel $\mathcal{C}''_1 : (bird \wedge \neg king.peng) \rightarrow flies$, where $\mathcal{C}_1 \sqsubseteq \mathcal{C}''_1$.[64] The other important property of this hierarchy is that the second proviso in the definition above is required to ensure that $\neg king.peng$ is not inserted as a background condition to the constraint $penguin \rightarrow \neg flies$. The reason for this is the following. There is a channel $(\mathcal{C}_4 ; \mathcal{C}_2 ; \mathcal{C}_1) \in \mathcal{T}(\mathcal{S})$ containing the constraint $king.peng \rightarrow flies$. Since $flies \perp \neg flies$, then $king.peng$ seems *prima facie* to be a background assumption to $penguin \rightarrow \neg flies$. However, the possibility of a regularity between $king.peng$ and $flies$ arises only because of the regularities between $king.peng$ and $penguin$, and between $penguin$ and $bird$—i.e. $\neg king.peng$ cannot be an implicit assumption underlying the constraint $penguin \rightarrow \neg flies$ because it is *only* via the property whereby king penguins are penguins that the conflicting regularity $king.peng \rightarrow flies$ arises. This is exactly the sort of case that the second opt-out clause in the definition is meant to cover. The kernel default hierarchy obtained from $\mathcal{S}$ is illustrated in Figure 4.1.

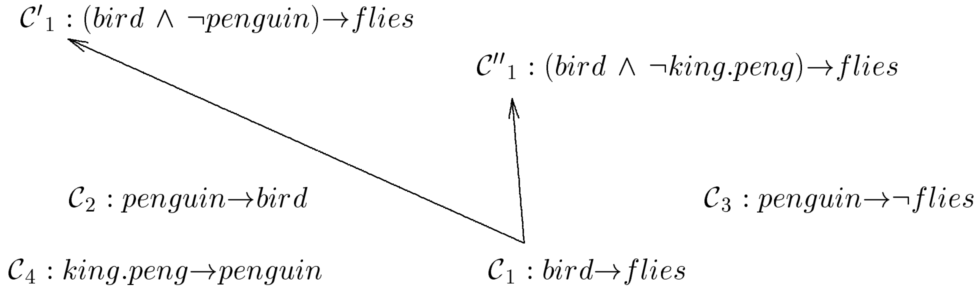

Figure 4.1: Channel Hierarchy for King Penguin Example

4.3.4 Inferring Generics

In the previous section, I defined the concept of a kernel default channel hierarchy, which effectively captures the background conditions associated with a given collection of channels. A kernel hierarchy is extended to a full default hierarchy by closing it

[64] It could be argued that the previous channel hierarchy already contains $\neg king.peng$ as an implicit assumption to the constraint $bird \rightarrow flies$ since $\neg penguin$ is a background assumption and every $king.peng$ token must also be a $penguin$ token. However, this view requires the channel-operations of Chapter 3 to be modified so as to obtain a suitable logic for generics, and my intention is to not modify them *at all*. Hence, *all* implicit background conditions need to be explicitly represented at some place in the hierarchy if they are to have an effect on the permitted patterns of inference.

under the channel operations, while ensuring that the following constraints of Section 3.5 are satisfied: Antecedent Background and Consequent Background. Recall that these constraints ensure that the representation of background conditions are symmetric between antecedent and consequent of channel-types.[65]

Definition *Let $\mathcal{S}$ be a collection of channels and $\mathcal{K}$ the corresponding kernel default hierarchy. The* full default channel hierarchy based on $\mathcal{S}$ *is the smallest channel hierarchy $\mathcal{H}$ containing $\mathcal{K}$ such that:*

1. *$\mathcal{H}$ is closed under the conditional channel operations;*
2. *$\mathcal{H}$ satisfies the Antecedent Background and Consequent Background Constraints.*

Given this definition of a channel hierarchy, inference of new generics from an initial collection is basically via the channel operations that are used to define it. These are the conditional channel operations of Chapter 3, used without modification—i.e. the *only* difference between the logic of generics and the logic of conditionals is that the channel hierarchy for the former is of a particular specified form.

Definition *Given a collection of channels $\mathcal{S}$, the collection $\mathcal{S}'$ of channels* defeasibly obtained from $\mathcal{S}$ *consists of the channels in $\mathcal{H}$, where $\mathcal{H}$ is the default channel hierarchy based on $\mathcal{S}$. I write $\mathcal{S} \mid\!\!\sim \mathcal{C}$ for $\mathcal{C} \in \mathcal{S}'$.*

It is hopefully clear that the difference between a logic of generics and a logic of conditionals resides only in what is acceptable as a channel hierarchy. In particular, in Chapter 3, a logic of conditionals is defined only with respect to some channel hierarchy (which is assumed to satisfy certain conditions, including Antecedent and Consequent Background Conditions), whereas the logic of generics requires a hierarchy to have specific properties. So, unlike previous attempts at defining logics of generics from conditional logics, the channel theoretic logic defined here effectively involves no modification to the *logic* of conditionals on which it is based—there is only a shift in perspective in that a logic of generics involves the assumption that any implicit background conditions to the given regularities are themselves extracted from those regularities themselves. This is captured by the use of a default channel hierarchy based on some specified kernel.

To be able to reason about the (default) properties of *individuals* using constraints, an assumption of *maximal normality* is required; otherwise, there is nothing to constrain the appearance of pseudo-signals and exceptions in channels. For example, if I

[65]The role of the other constraints of Section 3.5 is briefly discussed later.

know that *tweety* is a bird and I have no other information regarding *tweety,* this still does not prevent *tweety* from being a possible pseudo-signal in the channel containing the constraint $bird \rightarrow flies$. Chapter 5 introduces a *maximal normality* property that underlies a system of defeasible reasoning.[66] Note that the inference of a new generic is nonmonotonic in the sense that we may have $\mathcal{S} \mathrel{|\!\!\leadsto} \mathcal{C}$ without necessarily having $\mathcal{S}' \mathrel{|\!\!\leadsto} \mathcal{C}$, even when $\mathcal{S} \subseteq \mathcal{S}'$.

4.3.5 Some Examples

At this point, I turn to some simple examples to illustrate properties of the channel theoretic logic of generics before returning to the general patterns of inference of Section 4.1.2. For each example, I present a diagram which displays the relevant part of the default channel hierarchy.

Penguin example

The first example I illustrate is the standard penguin example. This demonstrates how Specificity is obtained—i.e. a "direct" generic overrides one that is obtained via a serial composition. The initial set $\mathcal{S}$ of channels is as follows.

$\mathcal{C}_1 : bird \rightarrow flies$
$\mathcal{C}_2 : penguin \rightarrow bird$
$\mathcal{C}_3 : penguin \rightarrow \neg flies.$

The kernel default channel hierarchy for $\mathcal{S}$ is depicted in Figure 4.2.

Some of the more interesting channels and constraints obtained by applying the (conditional) channel operations to the channel hierarchy include the following:[67]

$\mathcal{C}_3^* : flies \rightarrow \neg penguin$
$(\mathcal{C}_1 \,;\, \mathcal{C}_3^*) : bird \rightarrow \neg penguin$
$\mathcal{C}_1^* : \neg flies \rightarrow \neg bird.$

These derived channels and constraints correspond to the following generic sentences.

[66] An important difference between Morreau's (1992a) process of normalisation and the maximal normality assumption of the following chapter is that the process of normalisation is also used in Morreau's model-theoretic definition of the truth of generics and not just in the inference of (default) properties of individuals. This difference is because of Morreau's reductionist account of generics, reducing the semantics of generics to quantification over normal individuals.

[67] Recall that $\mathcal{C}_3^*$ is the contraposition channel of $\mathcal{C}_3$.

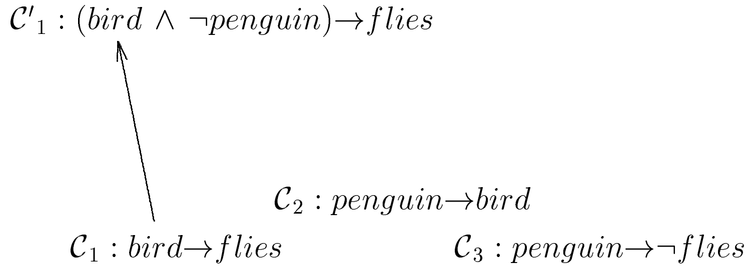

Figure 4.2: Channel Hierarchy for Simple Penguin Example

> "Things that fly (typically) aren't penguins"
> "Birds (typically) aren't penguins"
> "Things that don't fly (typically) aren't birds".

(I should mention my use of the word "typically" in the above generics. If this word is omitted, particularly from the third sentence above, then it is not so obvious that the resulting generics do indeed follow from those originally given. However, I am trying to capture a rather strong notion of consequence for generics (matching that considered acceptable by the AI literature); in particular, all the above are deemed consequences by Lifschitz (1989). Inserting "typically" makes the inference more reasonable and avoids the particular linguistic quirks with which I am not concerned.)

Non-taxonomic penguins

The previous example actually demonstrates the validity of what Morreau calls the *Taxonomic* Penguin Principle. This is an instance of Specificity in which one of the regularities that makes up the composite (overridden) regularity is taxonomic, or strict—i.e. the general rule cannot have any exceptions. For instance, *every* penguin is a bird—there can be no exception to this rule. A standard example from the literature that illustrates the standard Penguin Principle, in which all regularities are defeasible, is the following:

> "Students are adults."
> "Adults are employed."
> "Students are not employed."

It is generally agreed that the third generic here overrides the conclusion (that students are employed) obtained by transitively composing the first two. It should be clear that the channel theoretic logic handles the general Penguin Principle in exactly the same way as it handles the Taxonomic Penguin Principle—i.e. there is nothing in the treatment of the *penguin* example above that relies on the $penguin \rightarrow bird$ regularity being exception-free. Morreau's (1992a) logic of commonsense entailment, on the other hand, does not handle this example satisfactorily—the transitive conclusion is not overridden. Morreau deals with this problem by extending his normalisation process so that it is *prioritised*—some predicates are normalised in preference to others. This is reminiscent of the approach to Specificity taken within the traditional AI frameworks for default reasoning, and is very much an *ad hoc* addition of the sort that Morreau disparages.

Some of the generics inferred from the above three using the channel theoretic system are the following:

> "Employed people generally aren't students."
> "Adults generally aren't students."
> "People who aren't employed generally aren't adults."

Nixon diamond

A second important class of example is the classic Nixon Diamond.

> "Quakers are normally pacifists"
> "Republicans are normally not pacifists".

These sentences give rise to the following channels:

> $\mathcal{C}_1 : quaker \rightarrow pacifist$
> $\mathcal{C}_2 : republican \rightarrow \neg pacifist.$

The pertinent part of the associated default channel hierarchy is illustrated in Figure 4.3.

From this hierarchy, the following two generics are obtained:

> "Quakers that are not Republicans are normally pacifists"
> "Republicans that are not Quakers are normally not pacifists".

These are the only generics of any interest that are obtained. Perhaps surprisingly, no channel supporting a generic of the following form is obtained:

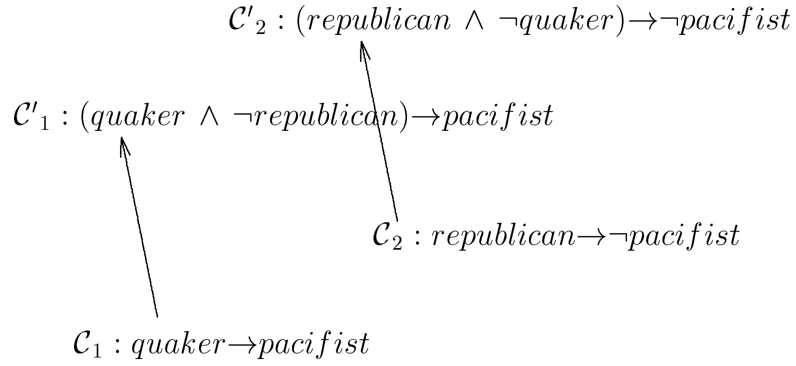

Figure 4.3: Channel Hierarchy for Nixon Diamond Example

> "Quakers are normally not Republicans"

(or the converse) since the background conditions of $\mathcal{C}_1$ conflict with those of $\mathcal{C}_2^*$. My intuitions aren't completely clear on this one; however, Lifschitz (1989) does not list this amongst the required inferences.

King Penguins

Finally, consider the following extension of the Penguin example, illustrating the use of the two provisos in the definition of default channel hierarchy (this example is isomorphic to the usual "Royal Elephant" example of the default inheritance literature).

The generics to be modelled are the following:

> "King penguins are penguins."
> "Penguins are birds."
> "Birds fly."
> "Penguins don't fly."

These are represented via the following set $\mathcal{S}$ of channels:

$\mathcal{C}_1 : king.peng \rightarrow penguin$
$\mathcal{C}_2 : penguin \rightarrow bird$
$\mathcal{C}_3 : bird \rightarrow flies$
$\mathcal{C}_4 : penguin \rightarrow \neg flies.$

The default channel hierarchy associated with $\mathcal{S}$ is shown in Figure 4.4.

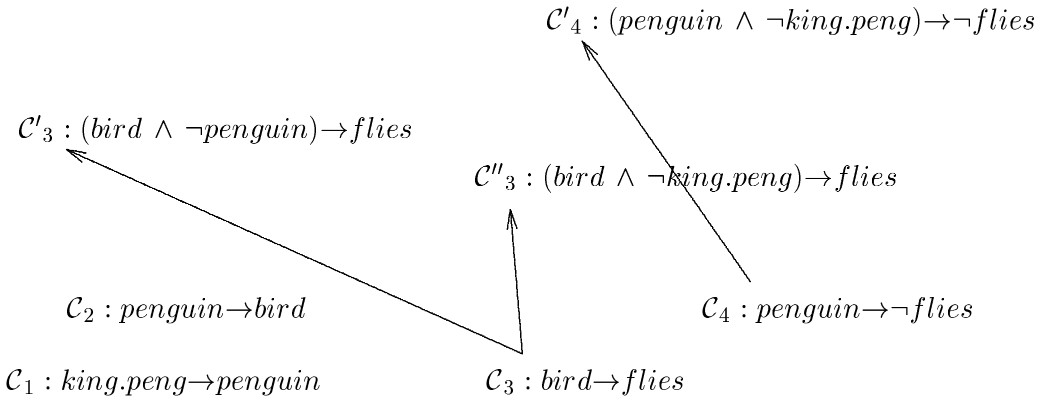

Figure 4.4: Channel Hierarchy for King Penguin Example

The channel $\mathcal{C}''_3$ is inserted into the above hierarchy because a channel containing $king.peng \rightarrow \neg flies$ is contained in $\mathcal{T}(\mathcal{S})$. There is also a channel in $\mathcal{T}(\mathcal{S})$ containing the constraint $king.peng \rightarrow flies$. However, this channel is the serial composition of a number of channels, one of which is $\mathcal{C}_2$, so by the definition of default channel hierarchy there is no channel in the hierarchy playing the same role as $\mathcal{C}''_3$ for $\mathcal{C}_4$—- i.e. $\neg king.peng$ is not inserted as a background condition to $\mathcal{C}_4$. From the hierarchy, $(\mathcal{C}_1 \,;\, \mathcal{C}_4) : king.peng \rightarrow \neg flies$ is obtained by serial composition. Note that if $\neg king.peng$ had been inserted as a background condition to $\mathcal{C}_4$, then this generic would not have been inferred.

4.3.6 Patterns of Inference Revisited

In this section, I show that the desirable patterns of inference described in Section 4.1.2 are supported by the channel theoretic logic of generics. For a comparison to be made between the channel theoretic logic and more standard logics, I need to ignore the role played by the supporting channels, of course, as was done with the logic of conditionals of Chapter 3. Given the discussion of Section 4.3.2, this is done by (i) assuming that all channels (reflexively) link the same classification and (ii) existentially quantifying over channels. This second point can be made more precise as follows.[68]

[68] The concept of the channel theoretic logic of generics supporting a pattern of inference is pretty much equivalent to the concept of the channel theoretic logic of conditionals supporting a rule of inference (Section 3.6.3), given the definition of '$\mid\!\leadsto$'.

Definition *The channel theoretic logic of generics validates the pattern of inference*

$$\frac{\begin{array}{c} A_1 \rightarrow B_1 \\ \ldots \\ A_n \rightarrow B_n \end{array}}{A \rightarrow B}$$

if for channels $\mathcal{C}_1, ..., \mathcal{C}_n$ *containing constraints*[69] $A_1^* \rightarrow B_1^*, ..., A_n^* \rightarrow B_n^*$ *respectively there is a channel* $\mathcal{C}$ *such that* $\{\mathcal{C}_1, ..., \mathcal{C}_n\} \mid\!\leadsto \mathcal{C}$ *and* $\mathcal{C}$ *contains the constraint* $A^* \rightarrow B^*$.

It will be convenient to extend the notation involving '$\mid\!\leadsto$' to involve constraints, rather than channels. In particular, I will sometimes write $\{\phi_1 \rightarrow \psi_1, ..., \phi_n \rightarrow \psi_n\} \mid\!\leadsto \phi \rightarrow \psi$ as an abbreviation for the following: there exist channels $\mathcal{C}_1, ..., \mathcal{C}_n, \mathcal{C}$ such that $\{\mathcal{C}_1, ..., \mathcal{C}_n\} \mid\!\leadsto \mathcal{C}$ and $\phi_1 \rightarrow \psi_1 \in typ(\mathcal{C}_1), ..., \phi_n \rightarrow \psi_n \in typ(\mathcal{C}_n), \phi \rightarrow \psi \in typ(\mathcal{C})$.

Given the above interpretation of the patterns of inference as they occur in standard logics of generics, I now turn to the patterns summarised in Section 4.1.2.

Specificity

I showed in Section 4.3.5 how the standard penguin example is handled. In general, $\{\phi \rightarrow \psi, \psi \rightarrow \tau\} \mid\!\leadsto \phi \rightarrow \tau$, but $\{\phi \rightarrow \psi, \psi \rightarrow \tau, \phi \rightarrow \neg\tau\} \not\mid\!\leadsto \phi \rightarrow \tau$.[70] It is easily seen why this is so—$\neg\tau$ conflicts with the consequent of $\phi \rightarrow \tau$, so $\neg\phi$ is inserted as a background assumption of $\phi \rightarrow \tau$, thereby preventing the conditional serial composition of $\mathcal{C}_1$ and $\mathcal{C}_2$ from containing the constraint $\phi \rightarrow \tau$.

Irrelevant conditions

The problem of Irrelevance is one that proves thorny to accounts of default reasoning based on logics of conditionals. Morreau's (1992a) process of normalisation ensures that he avoids this particular problem. The problem does not arise for the channel theoretic logic because of the assumption that all background conditions to the given constraints are contained within the constraints themselves. That is, any condition that is "irrelevant" (in the sense that it is not mentioned explicitly in the given set of constraints) will not be part of the background conditions of those constraints and

[69] Recall that the *-operator (introduced in Section 3.6.3) maps sentences of some propositional language to corresponding types.

[70] This also shows that Defeasible Transitivity is supported by the channel theoretic model.

will therefore have no effect on the channel operations or on the patterns of inference supported. So, for example:

$$\{bird{\rightarrow}flies\} \mid\!\sim (bird \wedge red){\rightarrow}flies.$$

However, if there happened to be a regularity between an object being red and an inability of it to fly, then this inference would no longer hold; i.e.

$$\{bird{\rightarrow}flies, red{\rightarrow}\neg flies\} \not\mid\!\sim (bird \wedge red){\rightarrow}flies$$

since in this case, $\neg red$ would be inserted as a background condition to $bird{\rightarrow}flies$ in the default channel hierarchy.[71]

Graded Normality and Defeasible Transitivity

Graded Normality is an important requirement for a logic of generics. In the channel theoretic system, this property is easily obtained since an exception in one channel need not be an exception in any other channel. Hence, even if *Leo* is an albino lion (e.g. $Leo \mapsto Leo$ is an exception in the channel $\mathcal{C}_1$ containing $lion{\rightarrow}brown$), he need not be an exception to the generic asserting he is dangerous (e.g. $(Leo \mapsto Leo : lion{\rightarrow}dangerous)$ may hold in channel $\mathcal{C}_2$).

The interaction of Graded Normality with Defeasible Transitivity is also not a problem for the channel theoretic framework. Given $\mathcal{S} = \{\phi{\rightarrow}\psi, \psi{\rightarrow}\tau, \psi{\rightarrow}\sigma, \phi{\rightarrow}\neg\sigma\}$, $\neg\phi$ is inserted as a background assumption to $\psi{\rightarrow}\sigma$ but not as a background assumption to $\psi{\rightarrow}\tau$. Hence, the serial composition of $\phi{\rightarrow}\psi$ with $\psi{\rightarrow}\tau$ is permitted, while the conditional serial composition of $\phi{\rightarrow}\psi$ with $\psi{\rightarrow}\sigma$ is blocked. That is, $\mathcal{S} \mid\!\sim \phi{\rightarrow}\tau$ but $\mathcal{S} \not\mid\!\sim \phi{\rightarrow}\sigma$.

Other patterns of inference

So far, I have examined several important patterns of inference from the literature on generics and default reasoning to show that the channel theoretic logic of generics

[71] This example illustrates the use of Defeasible Monotonicity. In general, $\{\phi{\rightarrow}\psi\} \mid\!\sim (\psi \wedge \tau){\rightarrow}\psi$, but $\{\phi{\rightarrow}\psi, (\phi \wedge \tau){\rightarrow}\neg\psi\} \not\mid\!\sim (\phi \wedge \tau){\rightarrow}\psi$. The second pattern of inference does not hold because $\neg(\phi \wedge \tau)$ is inserted as a background assumption to $\phi{\rightarrow}\psi$ in the default channel hierarchy, preventing the conditional serial composition with the constraint $(\phi \wedge \tau){\rightarrow}\phi$ from the logical channel. Note that by the definition, $\neg\phi$ is *not* inserted as a background condition to $(\phi \wedge \tau){\rightarrow}\neg\psi$ because of the presence of $(\phi \wedge \tau){\rightarrow}\phi$ in the logical channel (i.e. the first proviso in the definition in Section 4.3.3 covers this case).

satisfies many important properties, as far as its inferential behaviour is concerned. In Chapter 3, a wide variety of axioms and rules of inference from standard conditional logics were examined as a way of measuring the power of the channel theoretic logic of conditionals. Not all of these are universally accepted of course, and I discussed some of the arguments in favour of rejecting those that were not supported by the channel theoretic logic.

In the default reasoning literature, there are a number of important pieces of work that discuss patterns of reasoning that a nonmonotonic consequence relation should satisfy (e.g. (Gabbay 1985; Makinson 1989; Kraus et al. 1990; Kraus et al. 1990; Lehmann 1989)). These patterns include some of those that are considered unacceptable in a logic of conditionals (e.g. *Rational Monotonicity* (Kraus et al. 1990) is effectively equivalent to the axiom *CV,* which Nute (1980) rejects as a suitable axiom in a logic of conditionals). These lower bounds on nonmonotonic consequence relations provide an important set of tests to which the channel theoretic logic of generics can be put. I will describe them in detail, and analyse the channel theoretic logic in light of them, after defining a logic of default reasoning in Chapter 5.

4.3.7 Further Examples

In this section, I present further examples illustrating the patterns of inferences supported by the channel theoretic logic of generics. Some of these examples are taken from collections of "benchmark problems" of default reasoning (e.g. (Delgrande 1992; Lifschitz 1989)), although I am still only concerned with the inference of new default rules at this point.

Multiple inheritance

This example, taken from (Lifschitz 1989), is an extension of the Nixon diamond example.

"Quakers are normally pacifists"
"Republicans are normally hawks"
"Pacifists are normally politically active"
"Hawks are normally politically active"
"Pacifists are not hawks".

From these generics[72] the following channels are obtained, as well as the channel hierarchy partially shown in Figure 4.5.

$\mathcal{C}_1 : quaker \rightarrow pacifist$

$\mathcal{C}_2 : republican \rightarrow hawk$

$\mathcal{C}_3 : pacifist \rightarrow active$

$\mathcal{C}_4 : hawk \rightarrow active$

$\mathcal{C}_5 : pacifist \rightarrow \neg hawk.$

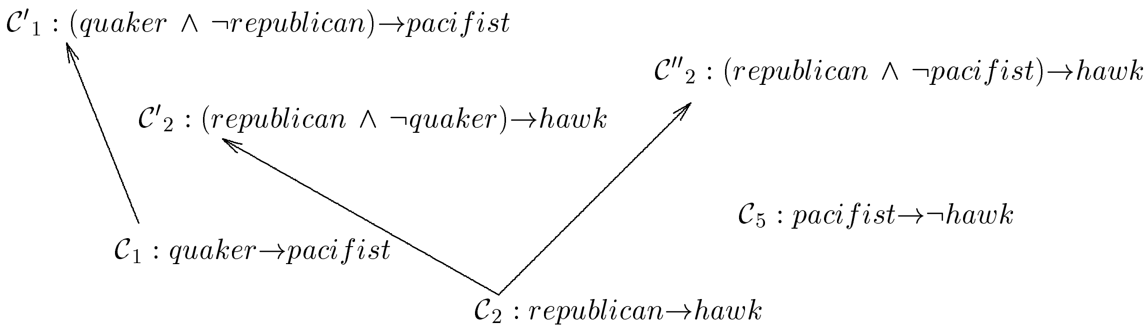

Figure 4.5: Partial Channel Hierarchy for Multiple Inheritance Example

The more interesting channels that can be constructed from the original ones are the following:

$(\mathcal{C}_1 \,;\, \mathcal{C}_3) : quaker \rightarrow active$

$(\mathcal{C}_2 \,;\, \mathcal{C}_4) : republican \rightarrow active$

$\mathcal{C}''_2 : (republican \wedge \neg quaker) \rightarrow hawk.$

The interesting point of this example is that we do not obtain the analogous constraint to the last one above, with the roles of *republican* and *quaker* reversed—i.e. there is no channel in the hierarchy supporting the constraint $(quaker \wedge \neg republican) \rightarrow hawk$.[73] The reason for this is that the contrapositive channel to $\mathcal{C}_5$ is not used when constructing the hierarchy (under the current definition of a default channel hierarchy)—it is the presence of $\mathcal{C}_5$ that results in channel $\mathcal{C}''_2$ being added to the default channel hierarchy (by definition), but without the contraposition of $\mathcal{C}_5$, we do not obtain an analogous

[72]Actually, the last of these is meant to be a *strict* relationship—i.e. one with no exceptions. I will treat it as a generic for now, but will comment on this point later.

[73]This conclusion is the only one required by Lifschitz that the system does not obtain.

background channel for $\mathcal{C}_1$, which is what is required for inferring the missing generic. The use of Contraposition (and Parallel Composition) in the construction of a default channel hierarchy is discussed in Section 4.4.3. Note, however, that the regularity supported by $\mathcal{C}_5$ is actually meant to be a "strict" one (i.e. it is not defeasible). If this regularity had been represented within the logical channel (as discussed in Section 4.3.2) then the contraposition of the regularity would have been available in the construction (since logical channels are closed under all the (standard) channel operations) and the above problem would have been averted.

King penguins revisited

Consider again the example involving king penguins from Section 4.3.5, and suppose the following channel is added to the initial set $\mathcal{S}$ of channels:

$\mathcal{C}_5 : \mathit{king.peng}{\rightarrow}\mathit{bird}.$

In this case, there *is* a channel $\mathcal{C} \in \mathcal{T}(\mathcal{S})$ containing the constraint $\mathit{king.peng}{\rightarrow}\mathit{flies}$ such that $\mathcal{C}$ that is not formed from a composition involving $\mathcal{C}_2$ ($\mathcal{C}$ is formed from the (standard) serial composition of $\mathcal{C}_5$ and $\mathcal{C}_3$). Hence, the resulting default channel hierarchy *does* contain a channel $\mathcal{C}'_4 : (\mathit{penguin} \wedge \neg \mathit{king.peng}){\rightarrow}\neg \mathit{flies}$, such that $\mathcal{C}_4 \sqsubseteq \mathcal{C}'_4$. The presence of this channel prevents the conditional serial composition of $\mathcal{C}_1$ with $\mathcal{C}_4$, and the desired generic

"King penguins don't fly."

is not obtained.

In standard inheritance reasoning, the link from *king.peng* to *bird* is considered to be *redundant*—any information that can be inferred along it can also be inferred along the composition of the links corresponding to $\mathcal{C}_1$ and $\mathcal{C}_2$. One could also argue that the link $\mathcal{C}_5$ is effectively redundant within the channel theoretic framework—e.g. it could be argued that $\mathcal{C}_5$ *is* just the serial composition of $\mathcal{C}_1$ and $\mathcal{C}_2$—and $\mathcal{C}_5$ being asserted as a channel in its own right does not prevent it being the case that $\mathcal{C}_5 = (\mathcal{C}_1 \,;\mathcal{C}_2)$. I believe that this argument carries some weight and would resolve the problem since the second proviso in the definition of default channel hierarchy would then apply. However, I will not press this point any further here.

4.4 Extensions to the Logic

In this section, I consider some possible extensions to the logic of generics defined above. Note that the logical machinery itself is not modified by any of these extensions—i.e. the channel operations remain equivalent to those defined in Chapter 3. The discussion of this section centres around increasing the representational power of the logic (e.g. representing strict, or taxonomic, relationships in a more natural way) and constraining the definition of a default channel hierarchy so that a more satisfactory logic results.

4.4.1 Representing the Dyadic GEN Operator

The dyadic *GEN* operator was introduced to account for certain generics with multiple readings. For example,

"Mary smokes after dinner."

has (at least) two readings, which can be represented respectively as follows:

$$GEN[x,s;](x = mary \wedge smoking(x,s) \;;\; after.dinner(s))$$
$$GEN[x,s;](x = mary \wedge after.dinner(s) \wedge in(x,s) \;;\; smoking(x,s)).$$

These two readings effectively involve interchanging the antecedent and consequent and the distinction between the two is thereby easily accounted for in the channel theoretic analysis.[74]

More generally, however, *GEN* involves parameters or variables which are not bound by the generic quantifier, as discussed in Section 4.1.3. For example, consider the following generic, taken from *TGB*:

"Typhoons arise over this part of the Pacific."

Again, this generic has (at least) two readings which can be represented as follows:

$$GEN[x;y](typhoon(x) \;;\; this.part.of.the.Pacific(y) \wedge arose.in(x,y))$$
$$GEN[x;y](this.part.of.the.Pacific(x) \;;\; typhoon(y) \wedge arose.in(y,x)).$$

In each of the above representations, the variable y is existentially quantified—e.g. in the first, given a (normal/typical) typhoon x, there is required to be *some* y such that

[74]The definition of situation theoretic objects which adequately represent the types described by the restrictor and matrix of each of these semantic forms is outside the scope of this thesis. However see (Cavedon and Glasbey 1994).

y is in "this part of the Pacific" and x arose there. We would like to be able to model this behaviour within the channel theoretic framework. One possibility is to use the framework described in Section A.3, incorporating situation theoretic parameters. In the treatment of parameters described in that section, any parameters which are contained only in the consequent of a constraint are effectively existentially quantified. For example, the first of the above generics involving typhoons could be seen as describing a channel containing a constraint of the following form, where all types are situation-types:

$$typhoon(X) \rightarrow this.part.of.the.Pacific(Y) \wedge arose.in(X, Y)$$

Given the treatment of parameters within the channel theoretic setting the generic describing this constraint could be paraphrased as follows:

> "In (normal) situations, for any typhoon there is somewhere in this part of the Pacific such that the typhoon arose there."

Under this interpretation, the generic quantification involves situations rather than typhoons—I am unclear as to whether this makes any significant difference linguistically.

A potential alternative approach, suggested to me by Robin Cooper, is to use n-ary types rather than parameters. Consider again the first reading of the above generic involving typhoons. This can be seen as describing a channel containing a constraint of the following form:

$$typhoon \rightarrow arose.in.this.part.of.the.Pacific.$$ [75]

The type *typhoon* is a unary type that classifies tokens as being typhoons while the type *arose.in.this.part.of.the.Pacific* is a binary type classifying a pair of tokens $\langle t, p \rangle$ exactly when p is in "this part of the Pacific" and t arose there. While this constraint does not involve any explicit quantification, its behaviour seems to involve implicit quantification of the sort required of the *GEN* operator. For example, consider a "normal" typhoon token t. If t is a normal token (with respect to the above constraint) then t is involved in a connection $t \mapsto \langle t, p \rangle$ that is classified by the constraint. But this would require there to be *some* "this part of the Pacific" token p such that t arose there—i.e. while t is generically quantified, p is effectively existentially quantified, which is the behaviour required for modelling the *GEN* operator.

[75] It is more general to view the consequent of this constraint as a conjunction of two types but it slightly simplifies the discussion to treat it as a single binary type.

The investigation of the use of n-ary types in constraints to obtain the quantificational behaviour required of the *GEN* operator is a topic for further research.

4.4.2 Strict Links

While an important tenet of channel theory is that any regularity can have exceptions, taxonomic relationships are exceptionless. Such relationships can be represented using the *type-entailment* relation. For example, "Penguins are birds" can be represented via the type entailment $bird \leq_A penguin$—this would ensure that any token in the classification A that was classified by the type *penguin* would also be classified by the type *bird*. Given the discussion in Section 4.3.2, such type-entailment relationships are also represented in any appropriate logical channel associated with a kernel default channel.[76] To allow a more consistent treatment of generics, however, exceptionless regularities are handled in the same way as those with exceptions. This requires the definition of a *strict channel*.[77]

Definition *A channel $\mathcal{C} : A \Rightarrow B$ is* strict, *written $\mathcal{C}!$, if none of the tokens of A are pseudo-signals and none of the tokens of $\mathcal{C}$ are exceptions to any of $\mathcal{C}$'s types.*

An important property that a logical channel possesses is that it effectively does not participate in a channel hierarchy—for a logical channel $\mathcal{L}$, if $\mathcal{L} \sqsubseteq \mathcal{C}$ then $\mathcal{L} = \mathcal{C}$. This condition is also required of strict channels.[78] Logical channels are also required to be closed under the channel operations. Since a strict channel does not have any background conditions encoded into it, then contraposing it or composing it with itself is equivalent under the standard and conditional operations. Further, the discussion of the following section ensures that contrapositions of strict channels are available when constructing a kernel hierarchy (alleviating the problem with the Multiple Inheritance

[76] In logical accounts of default reasoning, such taxonomic relationships are simply represented via a material implication. Default inheritance networks do not generally allow strict inheritance links (an exception is the work by Horty and Thomason (1988)). As should become clear below, if all taxonomic relationships were represented via type-entailment and logical channels, then the problem identified in the Multiple Inheritance example in Section 4.3.7 is avoided.

[77] In his early presentation of channel theory, Barwise (1993) defines the concept of a constraint $\phi \rightarrow \psi$ being ***absolute*** on a channel $\mathcal{C}$, meaning that $\phi \rightarrow \psi$ had no pseudo-signals with respect to $\mathcal{C}$. Since the concept of a connection being an exception in a channel does not exist in Barwise's formulation, absolute constraints are exactly those for which there are no counter-examples. Barwise uses the notation $c : \phi \overset{!}{\rightarrow} \psi$ to denote the fact that $\phi \rightarrow \psi$ is an absolute constraint.

[78] Actually, it need not be—since a strict channel $\mathcal{C}!$ contains no pseudo-signals or exceptions, any background assumptions encoded via a super-channel have no effect on $\mathcal{C}!$.

example of Section 4.3.7).

4.4.3 Tightening the Hierarchy—Contraposition and Parallel Composition

In the definition of a kernel default hierarchy, it was necessary to consider the (standard) transitive closure of channels when determining the defeaters of a constraint—i.e. when determining which types needed to be inserted as background conditions to a given constraint. This necessity was illustrated by the King Penguin example. However, it was noted that it was also necessary to close the channels under contraposition and parallel composition if an adequate logic is to be obtained. For example, consider the standard Penguin example but suppose that the "Penguins don't fly" generic is replaced by "Flying things aren't penguins":

$\mathcal{C}_1 : bird{\rightarrow}flies$
$\mathcal{C}_2 : penguin{\rightarrow}bird$
$\mathcal{C}_3 : flies{\rightarrow}\neg penguin.$

The kernel hierarchy for this example does not require any background conditions to be inserted for $\mathcal{C}_1$. When one then applies the conditional operators, $\mathcal{C}_1$ and $\mathcal{C}_2$ can be serially composed to yield a channel that contains the constraint $penguin{\rightarrow}flies$.

To ensure that $\neg penguin$ is inserted as a background condition to channel $\mathcal{C}_1$, the contraposition channel $\mathcal{C}_3^*$, which contains $penguin{\rightarrow}\neg flies$, needs to be involved in the construction of the kernel hierarchy. For similar reasons, parallel composition also needs to be involved, leading to the following revision of the definition of a kernel channel hierarchy.

Definition *Let $\mathcal{S}$ be a collection of channels containing a logical channel $\mathcal{L}$. Let $\mathcal{T}(\mathcal{S})$ be the collection of channels obtained by closing $\mathcal{S}$ under standard serial and parallel composition and standard contraposition. The* kernel default channel hierarchy for $\mathcal{S}$ *is the smallest hierarchy $\mathcal{K}$ of channels containing $\mathcal{S}$ such that:*

> *for each channel $\mathcal{C}_1 : \phi{\rightarrow}\psi$ in $\mathcal{S}$ and each channel $\mathcal{C}_2 : \tau{\rightarrow}\psi'$ in $\mathcal{T}(\mathcal{S})$, such that $\psi \perp \psi'$, $\mathcal{K}$ contains the channel $\mathcal{C}'_1 : (\phi \wedge \neg\tau){\rightarrow}\psi$, where $\mathcal{C}_1 \sqsubseteq_f \mathcal{C}'_1$ and $f(\phi{\rightarrow}\psi) = (\phi \wedge \neg\tau){\rightarrow}\psi$;* unless *either*
>
> 1. *there is a channel $\mathcal{C} : \phi{\rightarrow}\tau$ in $\mathcal{T}(\mathcal{S})$;* or
> 2. *$\mathcal{C}_2$ is formed by a composition involving a channel $\mathcal{C}' : \phi{\rightarrow}\sigma$, for some σ.*

This revised definition of a kernel hierarchy serves to make the logic *weaker*—since there are potentially more constraints (formed by parallel composition and contraposition) in $\mathcal{T}(\mathcal{S})$, then there are potentially more defeaters to each constraint. Hence, there are potentially more background conditions inserted into the hierarchy, and therefore potentially fewer new constraints that can be inferred via the conditional channel operations. As a specific illustration, the example above (containing $flies \rightarrow \neg penguin$) will now have $\neg penguin$ inserted as a background assumption to $\mathcal{C}_1$, preventing the unwanted transitive inference.

4.4.4 Strengthening the Logic by Constraining the Hierarchy

In Chapter 3, I showed how a channel hierarchy could be constrained to satisfy certain conditions, thereby ensuring that certain patterns of inference were supported by the channel theoretic logic of conditionals. So far, the only such constraints I have imposed on a default hierarchy is that it satisfies the Antecedent and Consequent Background Conditions. If the other conditions are also imposed, particularly Reliability and Consequent Consistency, then the logic of generics will be similarly strengthened. In fact, it turns out that these conditions are needed if the logic of default reasoning of Chapter 5 is to satisfy all of Kraus et al's (1990) required patterns of defeasible reasoning. This issue is discussed in Chapter 5.[79]

Of course, the logic of generics as it stands is dependent on the definition of a kernel default hierarchy—changing this definition would result in some different patterns of inference being supported while others would become invalidated. The current definition is a fairly simple one, designed to reflect the Fundamental Observation that the background conditions to a set of generics or default rules are somehow implicitly contained in that set. Revisions to the definition of a default hierarchy may well result in a logic of generics that better captures intuitively acceptable patterns of inference, but this is not of great concern to me here. The main contribution of the channel theoretic logic of generics is to show that the framework in which the logic of conditionals was constructed (in Chapter 3) can be easily modified to accommodate a logic suitable for generics *without* any modification to the "logical" operations themselves. The fact that

[79]Actually, I will take a rather different approach to ensuring that conditions similar to Reliability and Consequent Consistency are satisfied—i.e. I will instead impose conditions on the signalling relations of channels.

the resulting system satisfies the important test principles described above shows that it is a useful one for reasoning about generics.

4.5 Discussion

The simple channel theoretic analysis of generics provided in this chapter addresses some important issues that have proved problematic to other accounts. In particular, the context (i.e. channel) supporting the regularity corresponding to a given generic seems to play an important role. Most importantly (to the current enterprise) is that a logic for reasoning about generics was obtained without a need to reduce the truth of generics to relationships involving individuals. This logic was obtained directly from the channel theoretic conditional logic of the previous chapter by incorporating a simple restriction on appropriate channel hierarchies.

The fact that the truth of a generic does not involve individuals (in the current analysis) means that an extra mechanism—an assumption of *maximal normality*—is required before the system defined in this chapter can be used for defeasibly reasoning about the properties of individuals. This is the topic of the following chapter.

Chapter 5

A Channel–Theoretic Model of Default Reasoning

In the previous chapter, I showed how the channel-theoretic framework for reasoning about conditionals could be modified to obtain a system for reasoning about generics. Since a channel-type (with which the descriptive content of a generic is associated) is a primitive object, this treatment effectively takes generics to be simply true or simply false (with respect to some channel). Morreau (1992a) notes that such treatments do not directly lead to a logic for reasoning defeasibly about the properties of individuals using generics—an extra assumption of "normality" is required.

In this chapter, I define a condition that enforces an agent's expectations about a channel (in particular, the pseudo-signals and exceptions associated with a channel) to be "maximally normal"—i.e. an individual is not treated as an exception to the rule unless the given evidence requires that it be so. The definition of such a condition leads to a logic for default reasoning, obtained from the system for reasoning about generics of the previous chapter. Some extra machinery is required since default reasoning is linked to an agent's expectations or limited set of knowledge—i.e. the set of valid inferences licensed by default inference is dependent on an agent's epistemic state, not just on the way the world is. This requires a notion of *projection:* given some information (i.e. a collection of propositions) that an agent holds in some epistemic state, a set of channels is used to "project" that information so as to infer further information. This concept was defined in Section 2.3.6 (*S&B* (1993) define a similar concept).

5.1 Preliminaries

Before defining the formal concepts required for the channel-theoretic model of default reasoning, I briefly discuss some important conceptual issues.

5.1.1 Contextual Default Reasoning

In Section 5.2, a *maximal normality* condition is defined which, when imposed on a channel hierarchy,[1] results in a system for reasoning defeasibly about tokens and the types they satisfy. This condition is defined with respect to a given hierarchy—whether a token is "normal" or not in a channel is determined by comparing what is known about the token to the implicit background conditions of the channel.

The resulting model of default reasoning corresponds to a methodology of *contextual* or *situated* reasoning. The maximal normality definition ensures that for any channel $\mathcal{C}$ in which some token t is a pseudo-signal, there is some superchannel $\mathcal{C}'$ of $\mathcal{C}$ in which t is not a pseudo-signal.[2] This channel $\mathcal{C}'$ can be seen as the "most appropriate context" in which to reason about the token t, given the initial collection of information involving t. This corresponds to "reasoning at a level"—an agent does not use the whole hierarchy when reasoning but rather reasons with some particular channel at some appropriate level of the hierarchy.

The methodology of situated reasoning that suggests itself involves an agent "choosing the appropriate level" in a hierarchy and reasoning with the regularity at that level. The factors affecting the choice of appropriate level are not easily determined. One possible factor is the information available to the agent. For example, since moving up a hierarchy in general involves moving to a channel involving a richer collection of types, an appropriate level in a hierarchy is one in which the initially given information uses only types available at that level. These issues are discussed in Section 5.4 and are demonstrated in an approach to AI's *qualification problem.*

[1] Actually, the condition is imposed on a hierarchy of *links* which is obtained from a hierarchy of channels.

[2] This is not quite true for reflexive channels, since information known about t may conflict with a consequent of a constraint in $\mathcal{C}$.

5.1.2 The Need for Normality

Suppose there is a channel $\mathcal{C} : A \Rightarrow B$ supporting the regularity that birds fly—i.e. $typ(\mathcal{C})$ contains the constraint $bird \rightarrow flies$. Now further suppose that A supports the information that *tweety* is a bird—i.e. $(tweety : bird)$ holds in A. From Section 2.3, we know that for this information to carry the information that *tweety* flies via $\mathcal{C}$, there needs to be a connection $tweety \mapsto tweety$ in $tok(\mathcal{C})$. As discussed in Chapter 4, whether or not a generic is supported by $\mathcal{C}$ is independent of the properties of the associated individuals. In particular, there is no requirement imposed on *tweety* that it participate in any such connection in $tok(\mathcal{C})$, even if nothing more is known about *tweety* other than that $(tweety : bird)$ holds, even if it is known that *tweety* can fly! This is effectively an occurrence of the potential problem discussed earlier, whereby accounts of generics based on "simple truth" may not lead to a system for default reasoning about the properties of the associated individuals. This occurs because the fact that some given regularity holds—in this case, that birds fly—does not provide any information about the associated individuals.

The solution to this problem is via some normality assumption—i.e. an assumption that individuals are normal except in circumstances where they are known not to be so. This is analogous to processes in other logical treatments of generics, such as normalisation in Morreau's Commonsense Entailment or the circumscription of abnormality predicates. An important difference is that the normality assumption is not built into the logic for reasoning about generics themselves—the truth of generics is separate from the processes involved in *using* generics to reason defeasibly about individuals. In the channel-theoretic model, an exception to a rule corresponds either to an individual that does not participate in a connection or to a connection that does not behave "normally"—i.e. pseudo-signals and exceptions to channel-types. In the framework described below, I will draw a distinction between when an abnormal individual corresponds to a pseudo-signal and when it participates in an exception. This distinction is based on the definition of inference using *links* (see Section 2.3.2).

Reasoning defeasibly with generics generally involves reasoning with incomplete information. For example, if *tweety* is classified as being a bird, but it is unknown what sort of bird *tweety* is, then it is reasonable to conclude that *tweety* flies. However, if it is known that *tweety* is a penguin then this is no longer a reasonable inference. As such, the process of default inference cannot be based directly on channels since the inference from ($tweety$: $bird$) to ($tweety$: $flies$) may rely on a connection $tweety \mapsto tweety$ (in the channel containing $bird{\rightarrow}flies$) that is not there. Instead, the model of default reasoning defined below requires an agent to construct an information link $\mathcal{L}$ (based on the channel $\mathcal{C}$ supporting the "Birds fly" regularity) which is then used to draw the default conclusion. It is this link $\mathcal{L}$ that satisfies the maximal normality property—i.e. that there are as few pseudo-signals as possible. The normality condition is formally defined below after the process of inference using links is described.

5.1.3 An Intuitive Interpretation of Maximal Normality

The term "normality" has been used a number of different ways in this thesis. In Chapter 2, I discussed how Barwise and Seligman (1993) consider a connection being classified by a constraint as signifying that "normal conditions" are in place. In the previous chapter, I surveyed Boutilier and Morreau's logics of "normality". And in this chapter, I define a notion of "maximal normality" which is required for default reasoning about individuals. These concepts are related.

The logics of normality described in the previous chapter make use of normality orderings on worlds—i.e. worlds are ordered in increasing normality.[3] A world w is more normal than another w' if w satisfies more of the expectations of what normally holds. For example, if "Birds normally fly" is asserted, then (in the absence of other information) a world in which birds fly is deemed more normal than one in which they don't.

Normality in the sense used by *S&B* is related but slightly different. To *S&B*, "normality" reflects those conditions that are in place when a given connection is classified by a given constraint in some particular channel. The normal conditions in this sense are all those conditions that are required to be in place so that the type of information-flow licensed by the channel could successfully take place for that particular connection and constraint. Of course, it is impossible to list all such conditions so we must be satisfied with knowing they are in place when the information-flow successfully occurs. The "agent's expectations" comes into play by the way she carves the world up into channel-tokens and channel-types—if a channel contains a constraint γ and connection c, then the agent "expects" that conditions are normal, i.e. that the information-flow

[3] This is a slight simplification, of course, as is apparent from the discussion from Section 4.1.4. However, this abuse is acceptable in the current informal setting.

involving this connection and constraint takes place. The link corresponding to this channel would support the corresponding inference via the model of inference described in Section 2.3.6.

The notion of maximal normality defined below takes the channel-theoretic notion of normality and incorporates the usual default reasoning assumption that "things are as normal as possible". As will be seen, this effectively involves assuming that a given connection is classified by a constraint unless the given information—i.e. the information that makes up the premise of the default inference—conflicts with this assumption. Of course, the channel theoretic model of inference involves links rather than channels, so the assumption of maximal normality involves assuming that a link's signalling relation contains all relevant pairs except those that would lead to the inference of conflicting information.[4] The way in which this condition is defined is actually less circular than this informal description suggests and is based on a default hierarchy and the subchannel relation. However, the intuitions underlying the definition below can be thought of in the above terms.

5.1.4 Pseudo-Signals and Exceptions Again

Channel theory allows for two major sorts of errors: pseudo-signals and exceptions. Some of the conceptual differences between these two types of errors were discussed in Section 2.3.4. However, whether an error-token should be classed as an exception or a pseudo-signal is at best vague, and a counterexample could be best represented as an exception in one view of the world but as a pseudo-signal in another (equally valid) view. When dealing with inference, however, making a token a pseudo-signal is the only way in which to prevent that token from being involved in an inference (assuming that the token satisfies the antecedent of the constraint in question) since inference proceeds along a connection and hence an exception will license inference (even if the conclusion is an erroneous one). As such, tokens that are "abnormal" with respect to the agent's point of view—which is what is modelled by the maximal normality condition below—will need to be pseudo-signals. For example, if *tweety* is known to be a penguin, then *tweety* will be a pseudo-signal in the channel containing $bird \rightarrow flies$. The crucial point is that I want to model *reasonable inference*: if an agent knows that *tweety* is a penguin,

[4] I originally used a definition very similar to this in an early channel-theoretic framework for modelling Reiter's Default Logic (Cavedon 1993).

then the agent should not infer that *tweety* can fly.

A comprehensive model of situated default reasoning still has use for the concept of exceptions. The logic of defaults below models what an agent infers, based on certain information. For example, if *tweety* is a penguin then $\langle tweety, tweety \rangle$ will be contained in the maximally normal signalling relation associated with the link L supporting $penguin \rightarrow \neg flies$, leading to the inference that *tweety* does not fly. However, if *tweety* happens to have a jetpack strapped to his back, then he may indeed be able to fly. In this case, the connection involving *tweety*—i.e. the connection corresponding to the above signalling relation tuple—would be an exception in the channel corresponding to L and the inference that *tweety* is unable to fly would be an erroneous one. Since a channel hierarchy may be incomplete—i.e. it may be that not all possible background assumptions to a channel are represented in a given default hierarchy—some of the resulting default inferences may be incorrect, even when all information necessary for a correct inference is available.[5] That said, exceptions will be pretty much ignored for the rest of this chapter, which is concerned with "reasonable inference".

5.2 A Logic for Default Reasoning

In this section, I present the formal concepts of the model of default reasoning. Before doing so, however, I review and extend some of the concepts relating to inference via information links (see also Section 2.3.6).

As usual, I assume a number of simplifications in what follows. These are basically the same as those described in Section 4.3.2 and are used to "decontextualise" the system of default reasoning so as to allow direct comparison with more standard logics of default reasoning. In particular, all channels and links are assumed to be reflexive (i.e. link a classification to itself) and it is assumed that there is only a single (global) classification A with which all channels and links are concerned.

[5]That is, even if I know that *tweety* is wearing a jetpack, if I don't know that this will allow him to fly then I will infer that (being a penguin) he can't.

5.2.1 Inference via Information Links

In Section 2.3.6, I defined the process of inference, which is based on information links. Such information links are based on channels, in that the constraints and signalling relation of a link L are derived from the types (i.e. regularities) and connections of some channel $\mathcal{C}$. Given a link L and collection of propositions Ψ, the predicted inferences from Ψ supported by L are denoted by $Consq_L(\Psi)$. Inference via a link in this way is said to *induce* a classification B from Ψ; alternatively, $Consq_L(\Psi)$ is the *projection* of Ψ along L. These concepts are all defined in Section 2.3.6.

Given the notion of inference via information links, what remains to be provided is the property of the particular links that are used in a channel-theoretic model of default reasoning. These links are based on the channels of a default hierarchy, as defined in the previous chapter, but are required to satisfy a condition that enforces a condition of maximal normality, as defined below, on the signalling relations of those links. I need to define several concepts related to links before defining the framework for default reasoning. The first of these is a *sublink* relation, analogous to the concept of a subchannel relation (see Section 3.3.1).[6] One prior piece of notation I need is the following: for a tuple $\langle \phi, \psi \rangle$, $ante(\langle \phi, \psi \rangle) = \phi$ and $succ(\langle \phi, \psi \rangle) = \psi$.

Definition *Let* $L : A \rightrightarrows B$ *and* $L' : A' \rightrightarrows B'$ *be information links and let* f *be a function from* $L^\wedge$ *to* $L'^\wedge$[7]. L *is a* $\leq_{A',B'}$ –sublink of L' wrt f, *written* $L \preceq_{f,A',B'} L'$, *if*

1. $A \angle A'$, $B \angle B'$;
2. $L^\vee \subseteq L'^\vee$;
3. *for all* $\langle \phi, \psi \rangle \in L^\wedge$, $\phi \leq_{A'} ante(f(\langle \phi, \psi \rangle))$ *and* $succ(f(\langle \phi, \psi \rangle)) \leq_{B'} \psi$ *(where* $f(\langle \phi, \psi \rangle) \in L'^\wedge$ *).*[8]

The concept of a link hierarchy mirrors that of a channel hierarchy. Of course, there is no classification relation associated with a link so there are fewer conditions required of a pair of links related by sublink than of a pair of channels related by subchannel. The following result is a useful aside.

Proposition 5.2.1 *For any channels* $\mathcal{C}$ *and* $\mathcal{C}'$, *if* $\mathcal{C} \sqsubseteq_f \mathcal{C}'$ *then* $Link(\mathcal{C}) \preceq_f Link(\mathcal{C}')$.

---

[6] The definition of an information link, and the related terminology, was presented in Section 2.3.2.
[7] Recall that $L^\wedge$ is the indicating relation and $L^\vee$ the signalling relation of the link L.
[8] Note that this condition is exactly analogous to the last condition of the definition of subchannel.

Proof See Appendix D. □

To obtain a logic of default reasoning from the system for reasoning about generics, I actually need to derive a link hierarchy from a particular channel hierarchy—i.e. given a default hierarchy, an appropriate link hierarchy is obtained from it.

Definition *Given a channel hierarchy $\langle C, \mathcal{F}\rangle$ (where C is a collection of channels and $\mathcal{F}$ is a collection of functions) the* sublink hierarchy constructed from $\langle C, \mathcal{F}\rangle$ *is defined as follows:*

1. *the collection of links in the hierarchy is $\{Link(\mathcal{C}) \mid \mathcal{C} \in C\}$;*
2. *the collection of functions required for the sublink relation is as follows: for all channels $\mathcal{C}_1, \mathcal{C}_2$ in C such that $\mathcal{C}_1 \sqsubseteq_f \mathcal{C}_2$, f' is the function such that $f'(\langle\phi_1, \psi_1\rangle) = \langle\phi_2, \psi_2\rangle$ iff $\phi_1 \to \psi_1 \in typ(\mathcal{C}_1)$, $\phi_2 \to \psi_2 \in typ(\mathcal{C}_2)$ and $f(\phi_1 \to \psi_1) = \phi_2 \to \psi_2$.*
3. *$Link(\mathcal{C}_1) \preceq_{f'} Link(\mathcal{C}_2)$ iff $\mathcal{C}_1 \sqsubseteq_f \mathcal{C}_2$.*

It is easily shown that the above defines a valid hierarchy of information links, given a valid hierarchy of channels to begin with. Moving from a default channel hierarchy to a corresponding hierarchy of links is the first step in defining a system for default reasoning. The central step involves the imposition of a maximal normality condition on the signalling relations of these links, modelling the property of default reasoning whereby as many consistent conclusions as possible are drawn. This is discussed in the next section.

5.2.2 Maximal Normality

I have already discussed how an assumption of maximal normality—i.e. the assumption that individuals are as normal as possible—is required to extend the logic of generics of Chapter 4 to yield a system appropriate to defeasibly reasoning about the properties of individuals. In this section, I discuss some of the properties that a normality condition should intuitively possess, given the channel hierarchy framework for reasoning about generics. I also define a specific normality condition.

In the channel-theoretic model of generics, a channel-type (i.e. constraint) represents a rule which is generally reliable but which may contain exceptions. Some of the cases in which an exception occurs is captured by the way in which the channel $\mathcal{C}$ supporting the rule is related to other channels within a given channel hierarchy. Individuals that are exceptional with respect to the regularity correspond to either pseudo-signals or exceptions in $\mathcal{C}$. In the framework for default reasoning, I deal with links rather than channels and an individual for which a regularity is inapplicable (with respect to some given information) is a pseudo-signal in the link corresponding to the channel supporting that regularity. If the system for default reasoning is to infer as much as consistently possible about individuals, then it seems that pseudo-signals should occur only in the following sorts of cases:

- **Given information conflicts with information that would be inferred.** For example, in the standard *penguin* example, if one is told that *tweety* is a bird and that *tweety* doesn't fly, then *tweety* should be a pseudo-signal in the link whose indicating relation corresponds to the constraint $bird \rightarrow flies$.
- **Inheritance from separate premises would conflict.** In this case, we have two different links that potentially lead to conflicting conclusions. For example, if in the *penguin* example one is told that *tweety* is both a bird and a penguin, then *tweety* should be a pseudo-signal in one of the two links respectively supporting $bird \rightarrow flies$ and $penguin \rightarrow \neg flies$.

The first of these cases is taken care of by ensuring that a token is a pseudo-signal if its being otherwise would lead directly to a conclusion that conflicts with the given information. That is, given a link $L : A \stackrel{\rightarrow}{\rightarrow} A$ whose indicating relation supports the constraint $\phi \rightarrow \psi$, if $(t : \phi)$ and $(t : \neg\psi)$ both hold in A, then t should be a pseudo-signal in L, since otherwise L would support the inference of $(t : \psi)$, which conflicts with $(t : \neg\psi)$.

The second case above is blocked by considering the sublink hierarchy that is obtained from a default channel hierarchy. In a hierarchy of links (or channels), the background conditions of a particular link L are made explicit in some link L' of which L is a sublink—i.e. $L \preceq_f L'$. The constraints in L' that correspond to those in L—for example, suppose $f(\langle \phi, \psi \rangle) = \langle \phi', \psi' \rangle$—are more "refined" versions of those in L. In particular, $\langle \phi', \psi' \rangle$ may *explicitly* contain a type that is an *implicit* background assumption to $\langle \phi, \psi \rangle$. Any such background condition must be one that conflicts with the succedent of some other regularity, given the definition of a default channel hierarchy in Chapter 4. For example, the background assumptions to the constraint $bird \rightarrow flies$ are going to be types of the form $\neg\phi$ such that there is some constraint $\phi \rightarrow \neg flies$ in some other link. As such, a token t for which $(t : bird)$ holds should be a pseudo-signal to this constraint if

$(t : \phi)$ holds, where $\neg\phi$ is a background assumption to the constraint $bird \rightarrow flies$. Such background assumptions are determined by searching upwards in the link hierarchy.[9]

These intuitions lead to the following definition.

Definition *Let $L : A \rightrightarrows A$ be an information link (which is part of a link hierarchy) containing a constraint $\phi \rightarrow \psi$, Ψ a collection of propositions (formed from the tokens and types of A), and $t \in tok(A)$ a token such that $(t : \phi) \in \Psi$. The pair $\langle t, t \rangle$ is said to be* internally inappropriate to the signalling relation $L^{\vee}$ wrt Ψ *if either:*

1. *$(t : \psi') \in \Psi$ and $\psi \perp \psi'$, or*
2. *there exists a link L' such that $L \preceq_f L'$, where $f(\langle \phi, \psi \rangle) = \langle \phi', \psi' \rangle$, and $(t : \tau) \in \Psi$, where $\tau \perp \phi'$.*

Later, in Section 5.3.2, I also define the concept of *external* inappropriateness and take a connection to be inappropriate if it is either internally or externally inappropriate. I will sometimes abbreviate the way I refer to the above concepts by saying a tuple is inappropriate in a link L, rather than in the signalling relation of L. Also, I will refer to a tuple that is not inappropriate in L as being appropriate in L.

The concept of *maximal normality* is defined as follows. I have used a biconditional—i.e. *if and only if*—to reflect that the concept defined corresponds to *maximal* normality. This reflects the intuition that defeasible reasoning involves inferring *as much* information as is (consistently) possible.

Definition *A collection $\mathcal{L}$ of links arranged in a sublink hierarchy is said to be* maximally normal wrt Ψ *(a collection of propositions) if for each link $L : A \rightrightarrows A$ in $\mathcal{L}$ and each token $t \in tok(A)$, $\langle t, t \rangle$ is contained in the signalling relation $L^{\vee}$ if and only if L is a logical link*[10] *or $\langle t, t \rangle$ is appropriate to $L^{\vee}$.*

A pair $\langle t, t \rangle$ being inappropriate in a link L does not mean that t is a pseudo-signal in L since t may fail to satisfy the antecedent of any constraint supported by L. However, the converse clearly holds: if t is a pseudo-signal in L then $\langle t, t \rangle$ is inappropriate in L. Further, if $(t : \phi)$ holds for some $\langle \phi, \psi \rangle \in L^{\wedge}$, then t is a pseudo-signal in L if and only if $\langle t, t \rangle$ is inappropriate in L. (This last property plays an interesting role in Section

[9] Due to mutually conflicting constraints being handled via the channel hierarchy, the bias (i.e. towards "more specific" information) required to obtain the Specificity Principle is automatically supported.

[10] By logical link, I mean that L is a link corresponding to a logical channel. Notice that this means that logical links contain no pseudo-signals.

5.4, where I outline a methodology for reasoning defeasibly under the channel hierarchy framework.)

While a collection of links is required to be maximally normal for it to be used as the basis for default reasoning, there is no guarantee that the hierarchy of links obtained from a given default channel hierarchy is itself maximally normal, since there is no guarantee that the channels in the hierarchy satisfy any related condition. Hence, given an arbitrary collection of links, we require a maximally normal version of it.

Definition *Let $\mathcal{L}$ be a collection of links arranged in a sublink hierarchy and Ψ a collection of propositions. The* maximally normal version of $\mathcal{L}$ wrt Ψ *is a collection of links $\mathcal{L}'$ exactly like $\mathcal{L}$, except possibly in the signalling relations of the links, where $\mathcal{L}'$ is maximally normal wrt Ψ.*

It can be shown that if $\mathcal{L}$ is a given hierarchy of links constructed from a default channel hierarchy,[11] then the maximally normal version of $\mathcal{L}$ preserves the hierarchy, in that the sublink relation is mirrored in the way illustrated by the following result.

Proposition 5.2.2 *Let $\mathcal{D} = \langle \mathcal{L}, \mathcal{F} \rangle$ be a link hierarchy and let $\mathcal{D}' = \langle \mathcal{L}', \mathcal{F} \rangle$ be a maximally normal version of $\mathcal{D}$. Further, let $L_1, L_2 \in \mathcal{L}$, where $L_1 \preceq_f L_2$, and let $L'_1, L'_2 \in \mathcal{L}'$ be the links corresponding to L_1, L_2 respectively. Then $L'_1 \preceq_f L'_2$.*

Proof See Appendix D. $\Box$

As an example of the above definitions, consider the standard *penguin* example. The default channel hierarchy (ignoring the token-level) for this example is shown in Figure 4.2 in Chapter 4. The corresponding sublink hierarchy (again ignoring the signalling relation) is illustrated in Figure 5.1.

Consider a *bird* token *tweety* and suppose that $(tweety : bird)$ holds, and that is all. $\langle tweety, tweety \rangle$ is not inappropriate to L_1 and so $\langle tweety, tweety \rangle$ must be in the signalling relation of L_1 (assuming that the collection of links in the figure satisfies maximal normality). This leads to an inference of $(tweety : flies)$. (Note that $\langle tweety, tweety \rangle$ is not inappropriate in L_3 since $(tweety : penguin)$ does not hold.) Now suppose that $(tweety : \neg flies)$ holds. In this case, $\langle tweety, tweety \rangle$ *is* inappropri-

[11] The restriction to a default channel hierarchy is required as otherwise the result may not hold—for example, the definition of sub-link allows a link L to possibly contain a type that does not correspond to any type in a sublink L' of it, which may result in the maximally normal version of L' not being a sublink of L. This should be made clear on examination of the proof of the following result.

$L_1 : (bird \wedge \neg penguin) \rightarrow flies$

$L_2 : penguin \rightarrow bird$

$L_1 : bird \rightarrow flies$ $\quad\quad$ $L_3 : penguin \rightarrow \neg flies$

Figure 5.1: Link Hierarchy for Penguin Example

ate in L_1, and *tweety* is therefore a pseudo-signal in this link, preventing the inference to $(tweety : flies)$. Similarly, if instead $(tweety : penguin)$ holds, then $\langle tweety, tweety \rangle$ is inappropriate in L_1 by virtue of *penguin* conflicting with one of the background conditions of L_1 as represented in L_1'. In this case, $\langle tweety, tweety \rangle$ is not inappropriate in L_3 and is therefore contained in the signalling relation of that link, leading to the inference of $(tweety : \neg flies)$.

5.2.3 Default Inference Systems

I am now able to define the channel-theoretic system for defeasible reasoning. The basis of the system is the logic of generics of Chapter 4, which is used to obtain a default channel hierarchy from an initial collection of channels. Given such a default hierarchy, a corresponding hierarchy of links is extracted. This collection $\mathcal{L}$ of links is used as the basis for defeasibly reasoning about tokens (i.e. individuals)—given a collection Ψ of propositions, each link in $\mathcal{L}$ is required to satisfy maximal normality and the information projected from Ψ via $\mathcal{L}$ is the information defeasibly inferred.

Definition *Let Ψ be a collection of propositions and $\mathcal{S}$*[12] *a collection of channels with corresponding default channel hierarchy $\mathcal{H}$, and let $\mathcal{L}$ be the sublink hierarchy constructed from $\mathcal{H}$. The* default inference system based on $\mathcal{S}$ and Ψ *is the link hierarchy $\mathcal{D}$ that is the maximally normal version of $\mathcal{L}$ wrt Ψ.*

It can be quite easily shown that there is a unique default inference system based on a given collection of channels $\mathcal{S}$ and collection Ψ of propositions.[13] Defeasible inference

[12] As usual, I assume that $\mathcal{S}$ contains an appropriate logical channel.

[13] Without giving all the details, the outline of the proof is as follows. Under the assumption that

from the information supported by Ψ is performed by projecting Ψ via the maximally normal hierarchy $\mathcal{D}$.

Definition *Let Ψ be a collection of (Austinian) propositions and $\mathcal{S}$ a collection of channels each supporting a generic or default rule. The collection of* default consequences of Ψ wrt $\mathcal{S}$ *is given by $\{(t:\phi) \mid (t:\phi) \in Consq_{\mathcal{D}}(\Psi)$, where $\mathcal{D}$ is the default inference system based on $\mathcal{S}$ and $\Psi\}$.*
I use the notation $\Psi \mathrel{|\!\!\sim}_{\mathcal{S}} p$ to denote that p is a default consequence of Ψ wrt $\mathcal{S}$.

5.2.4 Some Examples

At this point, it is useful to illustrate the concepts defined so far by considering some simple examples. Note that most of the logical machinery in a default inference system resides in the logic of generics—i.e. the system by which new channels are obtained from the given ones by way of the conditional channel operations defined in Chapter 3. The maximal normality condition is simply used to ensure that each token appears in the signalling relations of the pertinent links when they should, ensuring that the corresponding individual is inferred (via projection) to possess some property.

Penguins

The standard *penguin* example was used to illustrate the channel theoretic logic of generics in Section 4.3.5. The collection of channels $\mathcal{S}$ given initially (i.e. corresponding to the given generics) contains the following:

$\mathcal{C}_1 : bird \rightarrow flies$
$\mathcal{C}_2 : penguin \rightarrow bird$
$\mathcal{C}_3 : penguin \rightarrow \neg flies.$

The default channel hierarchy (illustrated in Figure 4.2) also contains the following channels:

---

there is one unifying classification (which fixes the collection of tokens and types for the whole system), only the definition of a kernel default hierarchy (Section 4.3.3) allows for any non-determinism in the properties of a default system, since the token-level connections of the channels encoding the background conditions are unrestricted. However, the signalling relations of all links in a default system are completely determined by the maximal normality condition.

$\mathcal{C}_3^* : flies \rightarrow \neg penguin$
$(\mathcal{C}_1 \,;\, \mathcal{C}_3^*) : bird \rightarrow \neg penguin$
$\mathcal{C}_1^* : \neg flies \rightarrow \neg bird.$

The links in the sublink hierarchy corresponding to the default channel hierarchy contain the obvious indicating relations corresponding to the constraints of these channels. The default inference system must also satisfy maximal normality, with respect to some collection Ψ of propositions.

For the first instance of this example, suppose that we have a token *tweety* that is classified as being a *bird*—i.e. $(tweety : bird) \in \Psi$. The pair $\langle tweety, tweety \rangle$ is inappropriate only in the link $Link(\mathcal{C}_1^*)$—i.e. that whose indicating relation contains $\langle \neg flies, \neg bird \rangle$. Hence, *tweety* is not a pseudo-signal in the maximally normal version of any of the other links, which allows the following conclusions to be drawn: $(tweety : flies), (tweety : \neg penguin)$—i.e.:

$\{(tweety : bird)\} \mid\!\!\sim_{\mathcal{S}} (tweety : flies)$
$\{(tweety : bird)\} \mid\!\!\sim_{\mathcal{S}} (tweety : \neg penguin).$

Now suppose that $(tweety : \neg flies)$ also holds in Ψ. In this case, $\langle tweety, tweety \rangle$ *is* inappropriate in $Link(\mathcal{C}_1)$ (as well as in $Link(\mathcal{C}_3^*)$ and $Link(\mathcal{C}_1^*)$) so *tweety* is a pseudo-signal in the maximally normal version of this link and the conclusion $(tweety : flies)$ is not drawn. If the only given information is $(tweety : \neg flies)$, then $\langle tweety, tweety \rangle$ is no longer inappropriate in $Link(\mathcal{C}_1^*)$ and the following is obtained:

$\{(tweety : \neg flies)\} \mid\!\!\sim_{\mathcal{S}} (tweety : \neg bird).$

If the only information given initially is $(tweety : penguin)$, then $\langle tweety, tweety \rangle$ is inappropriate in $Link(\mathcal{C}_1)$ and the required conclusions are obtained:

$\{(tweety : penguin)\} \mid\!\!\sim_{\mathcal{S}} (tweety : bird)$
$\{(tweety : penguin)\} \mid\!\!\sim_{\mathcal{S}} (tweety : \neg flies).$

However, if $(tweety : flies)$ is added to the premise, then only the first of these is obtained.

Multiple Inheritance

The extended version of the standard *Nixon Diamond* example was described in Section 4.3.7. The collection $\mathcal{S}$ of given channels contains the following:

$\mathcal{C}_1 : quaker \rightarrow pacifist$

$\mathcal{C}_2 : republican \rightarrow hawk$

$\mathcal{C}_3 : pacifist \rightarrow active$

$\mathcal{C}_4 : hawk \rightarrow active$

$\mathcal{C}_5 : pacifist \rightarrow \neg hawk$

and the corresponding default hierarchy is illustrated in Figure 4.5. This hierarchy also contains the following channels:

$(\mathcal{C}_1 \,;\, \mathcal{C}_3) : quaker \rightarrow active$

$(\mathcal{C}_2 \,;\, \mathcal{C}_4) : republican \rightarrow active.$

Let *nixon* be a token regarding which the following information is known: $\{(nixon : quaker), (nixon : republican)\}$.[14] The pair $\langle nixon, nixon\rangle$ is inappropriate in both $Link(\mathcal{C}_1)$ and $Link(\mathcal{C}_2)$ (due to channels $\mathcal{C'}_1$ and $\mathcal{C'}_2$—see Figure 4.5) and therefore no conclusion is obtained as regards whether *nixon* is a *pacifist* or a *hawk*. However, $\langle nixon, nixon\rangle$ is inappropriate in neither of $Link(\mathcal{C}_1 \,;\, \mathcal{C}_3)$ nor $Link(\mathcal{C}_2 \,;\, \mathcal{C}_4)$ and hence the following conclusion is obtained:

$$\{(nixon : republican), (nixon : quaker)\} \mathrel{|\!\!\sim}_{\mathcal{S}} (nixon : active).$$

The interesting point to note from this example is that the reasoning corresponds to that termed *skeptical* (as opposed to *credulous*) in the inheritance networks literature (e.g. (Touretzky et al. 1987)). Skeptical reasoners are only able to draw conclusions that follow "unambiguously" in the sense that two distinct readings of the network cannot be obtained. For example, a credulous reasoner could be used to infer that *nixon* is a *pacifist*, or it could be used to infer that *nixon* is a *hawk*, but not both conclusions at once. Skeptical reasoning corresponds to allowing only conclusions that follow in *all* default extensions when using Reiter's Default Logic, or *all* stable autoepistemic extensions when using autoepistemic logic, or *all* models when using a model-theoretic nonmonotonic logic, such as circumscription or commonsense entailment. The channel-theoretic system reasons skeptically because of the way default hierarchies are constructed—since all potential background conditions are inserted when the succedents of constraints conflict then the pertinent pairs of the signalling relations will be inappropriate to all such

[14] I haven't been explicit about it before, but it should be noted that the incorporation of a logical channel ensures that the form in which information is presented does not affect the behaviour of the logical system. In particular, if the initial set of propositions is $\{(nixon : republican \wedge quaker)\}$ then the same conclusions are reached.

constraints.

Graded Normality

A final simple example serves to demonstrate that Graded Normality is supported at the level of individuals—i.e. an individual that is abnormal in one respect need not be abnormal in any other. As discussed in Section 4.2.2, this property is supported due to the fact that different generics are supported by different channels, and a token can be a pseudo-signal in one channel without it necessarily being a pseudo-signal in any other.

To illustrate, let $\mathcal{S}$ be a collection of channels containing the following:

$\mathcal{C}_1 : lion {\rightarrow} dangerous$
$\mathcal{C}_2 : lion {\rightarrow} brown$

and let Ψ be the following collection of propositions: $\{(leo : lion), (leo : \neg brown)\}$. The tuple $\langle leo, leo\rangle$ is inappropriate in $Link(\mathcal{C}_2)$ but not in $Link(\mathcal{C}_1)$, which means that leo is a pseudo-signal in the maximally normal counterpart of $Link(\mathcal{C}_2)$ but not in that of $Link(\mathcal{C}_1)$—i.e.:

$$\{(leo : lion), (leo : \neg brown)\} \mathrel{|\!\!\sim}_{\mathcal{S}} (leo : dangerous).$$

5.3 Evaluating the Logic of Defaults

In this section, I present some results which help to establish the patterns of reasoning supported by the channel theoretic logic of default reasoning.

5.3.1 Nonmonotonic Consequence Relations

Recent work in the AI literature has been concerned with the investigation of *nonmonotonic consequence relations*. In the spirit of Tarski's work defining the minimal conditions that a consequence relation for a (classical) logic should satisfy, several authors have defined general conditions that a nonmonotonic consequence relation should satisfy (e.g. (Gabbay 1985; Makinson 1989; Lehmann 1989; Kraus et al. 1990; Freund et al. 1991; Freund and Lehmann 1993)).[15] Gabbay (1985) argued that Tarski's con-

[15] Also of interest is Shoham's (1987) characterisation of nonmonotonic logics using the notion of *preferred models:* i.e. Shoham shows that different nonmonotonic logics can be characterised by defining

dition of *Monotonicity* should be dropped in favour of a condition he called *Cautious Monotonicity* (see below). Probably the most extensive investigation of nonmonotonic consequence relations and their properties has been performed by Lehmann and his colleagues (e.g. (Lehmann 1989; Kraus et al. 1990; Freund et al. 1991; Freund and Lehmann 1993)), and it is this work on which I will base the following discussion. (In the following, I will use *KLM* to refer to (Kraus et al. 1990) and *FLM* to refer to (Freund et al. 1991)). *KLM* and *FLM* give both syntactic and model-theoretic characterisations of the conditions that they require of nonmonotonic consequence relations, but I will discuss only the syntactic characterisations here since these are more readily utilised for analysing the properties of the channel-theoretic default logic.

KLM's syntactic characterisation of nonmonotonic consequence relations is via Gentzen-style sequents and the specification of minimal conditions that such relations must satisfy. A *conditional assertion* is a syntactic object of the form $A \mathrel{|\!\sim} B$, where A and B are sentences of some language and '$\mathrel{|\!\sim}$' denotes nonmonotonic consequence. The intended interpretation of such an object is "from A, one may sensibly conclude B". Minimal conditions on consequence relations are specified in the following definition. Note that a nonmonotonic consequence relation is assumed to be defined on top of some minimal monotonic consequence relation—monotonic consequence is denoted by '$\models$'. As usual, I use '$\supset$' to denote material implication and '$\Leftrightarrow$' to denote logical equivalence.

Definition *(KLM) A* preferential consequence relation *is characterised by the following conditions:*

Left Logical Equivalence (LLE) $$\frac{\models A \Leftrightarrow B \qquad A \mathrel{|\!\sim} C}{B \mathrel{|\!\sim} C}$$

Right Weakening (RW) $$\frac{\models A \supset B \qquad C \mathrel{|\!\sim} A}{C \mathrel{|\!\sim} B}$$

Reflexivity (Refl) $$A \mathrel{|\!\sim} A$$

And $$\frac{A \mathrel{|\!\sim} B \qquad A \mathrel{|\!\sim} C}{A \mathrel{|\!\sim} A \wedge C}$$

orderings between models and defining validity via truth in minimal models rather than in all models.

Or $$\frac{A \mid\sim C \quad B \mid\sim C}{A \vee B \mid\sim C}$$

Cautious Monotonicity (CM) $$\frac{A \mid\sim B \quad A \mid\sim C}{A \wedge B \mid\sim C}$$

Kraus et al. argue that each of the above conditions is acceptable for nonmonotonic consequence and illustrate this using various examples. They also show that preferential relations can be characterised using *preferential models*, which impose an ordering on possible assignments to predicates, similar to Shoham's (1987) notion of preferred model. *FLM* further investigates nonmonotonic consequence relations and in particular argues that nonmonotonic consequence should also satisfy the following condition.

Definition *(FLM) A* rational consequence relation *is a preferential consequence relation that satisfies the following extra condition:*

Rational Monotonicity (RM) $$\frac{A \mid\sim C \quad A \not\mid\sim \neg B}{A \wedge B \mid\sim C}$$

KLM contains several revealing results regarding rational relations and the *RM* condition in particular. For example, they introduce the following conditions and prove the theorem below.

Weak Transitivity (WT) $$\frac{A \mid\sim B \quad B \mid\sim C \quad B \not\mid\sim \neg A}{A \mid\sim C}$$

Weak Contraposition (WC) $$\frac{C \wedge A \mid\sim B \quad C \not\mid\sim B}{C \wedge \neg B \mid\sim \neg A}$$

Theorem 5.3.1 *(FLM) For preferential relations, Rational Monotonicity is equivalent to Weak Transitivity, and implies Weak Contraposition.*

This result is of some interest in the context of conditional reasoning. If, as *FLM* claim, rationality is a reasonable condition to require a consequence relation to satisfy, then the above result demonstrates why Contraposition and Transitivity are generally considered tenable patterns of inference: the "weak" versions of these rules *are* valid for rational relations, and are equivalent to the standard versions under fairly general conditions.

In the following sections, I will show how default consequence, defined with respect to channel-theoretic default systems, measures up against the *KLM* and *FLM* principles. These principles are not satisfied by all logics of default reasoning[16] but do provide a useful generic testbed. Note, however, that these principles are not all well suited to a logic of generics or default reasoning that can involve different contexts. For example, *And* would be an unreasonable pattern of inference if the premises were supported in different contexts. (This is one of my reasons for the simplifying assumption that all channels link the same classification to itself).

5.3.2 Strengthening Normality

In Section 3.5, a number of conditions were defined which, when satisfied by a channel hierarchy, guarantee that certain axioms and rules of inference of traditional conditional logics are supported by the channel-theoretic system defined in Chapter 3. I argued there that these conditions were reasonable ones to expect a channel hierarchy to satisfy, regardless of the logical properties consequently enjoyed by the framework.

The channel hierarchy conditions can be partitioned into two classes: the *symmetry* constraints, consisting of the Antecedent Background, Consequent Background, Parallel Subchannel and Serial Subchannel Constraints; and the *reliability* constraints, consisting of the Reliability and Consequent Consistency Constraints. This partition is based on the observation that the symmetry constraints are basically concerned with ensuring that a channel hierarchy is symmetric (in the way it encodes background assumptions) between antecedents and consequents and that a hierarchy interacts smoothly with the channel operations; while the reliability constraints are directed at ensuring that channels are not inherently unreliable in the manner demonstrated in Section 3.5. Another important difference is that a hierarchy that doesn't satisfy the symmetry constraints can be modified to one that does, without changing its fundamental properties, by adding certain channels if necessary. Based on this last property, I will henceforth assume that any default channel hierarchy satisfies the symmetry constraints.

[16] In particular, Zhang and Rounds (1993) argue that Cautious Monotonicty should not be validated by a logic of default reasoning. On the other hand, Boutilier (1992) has shown that his logic *CT4* defines a preferential consequence relation while his logic *CT4D* defines a rational consequence relation.

The Reliability and Consequent Consistency Constraints, as they are defined in Section 3.5, are concerned with the "internal" reliability of channels, ensuring that the types linked by constraints in a given channel are not such that the channel is guaranteed to contain exceptions. However, since the default logic defined here has been "decontextualised" so as to allow direct comparison with more standard principles of default reasoning, the reliability constraints need to be generalised so as to apply across multiple channels. For example, consider the Reliability Constraint, which requires two constraints with the same antecedent to have non-conflicting background assumptions. In Section 3.5.5, the Reliability Constraint was only required where such pairs of constraints were from the same channel. However, consider again the *peacock* example described earlier which involves (different) channels $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively supporting the following regularities:

$$peacock{\rightarrow}bright.feathers$$
$$peacock{\rightarrow}lays.eggs.$$

These regularities have conflicting background assumptions—i.e. the first involves the assumption that a peacock is male while the second involves the assumption that the peacock is female. As such, the (default) inference that a given peacock both has bright feathers and lays eggs should not be allowed.

Of course, I do not want to rule out channels of this form, which would be the case if Reliability was required of the corresponding hierarchy.[17] Rather, I wish to ensure that inference behaves properly—i.e. that the associated signalling relations of the maximal normal links do not lead to the inference of an incoherent collection of information. The current version of normality is not strong enough to ensure this for all cases. For example, in the *peacock* example, if all that is known about the individual *pat* is that it is a peacock—i.e. $\{(pat : peacock)\}$ is the premised information—then there is no conflict with the background assumptions of either of the above regularities. Hence, the tuple $\langle pat, pat \rangle$ is not inappropriate (by the definition in Section 5.2.2) in either of the links L_1 or L_2 supporting these regularities and therefore leads to the inference that *pat* both lays eggs and has bright feathers. Coupled with links corresponding to generics "Peacocks with bright feathers are male" and "Egg-laying peacocks are female", this

[17] Even if I were to impose Reliability and Consequent Consistency on hierarchies, then it would be much preferable to determine the conditions which a collection $\mathcal{S}$ of channels had to satisfy so as to ensure that the default hierarchy corresponding to $\mathcal{S}$ satisfied Reliability, rather than to simply require all default hierarchies to satisfy Reliability.

could further lead to the inference that *pat* is both male and female.

This problem is rectified by the following definition of *external inappropriateness.* This concept plays a similar role to Reliability and Consequent Consistency, but works across channels and involves the token level. Other than the fact that it involves more than one channel, the following concept is similar to that of internal inappropriateness defined previously.

Definition *Let* $L_1 : A \rightrightarrows A$ *and* $L_2 : A \rightrightarrows A$ *be information links (which are part of some link hierarchy) such that* $\langle \phi, \psi \rangle \in L_1^\wedge$ *and* $\langle \tau, \sigma \rangle \in L_2^\wedge$. *Further, let* Ψ *be a collection of propositions formed from the types and tokens of* A *and let* $t \in tok(A)$ *be such that* $(t : \phi)$ *holds in* Ψ *and either* $(t : \tau)$ *also holds in* Ψ *or* $\tau \leq_A \psi$. *The pair* $\langle t, t \rangle$ *is said to be* externally inappropriate to the signalling relations $L_1^\wedge$ and $L_2^\wedge$ wrt Ψ[18] *if:*

1. $\langle t, t \rangle$ *is internally inappropriate in neither* L_1 *nor* L_2 *wrt* Ψ*;*
2. *there exist links* L_1', L_2' *such that* $L_1 \preceq_{f_1} L_1'$, $L_2 \preceq_{f_2} L_2'$, *where* $f_1(\langle \phi, \psi \rangle) = \langle \phi', \psi' \rangle$ *and* $f_2(\langle \tau, \sigma \rangle) = \langle \tau', \sigma' \rangle$, *and either* $\phi' \perp \sigma$ *or* $\tau' \perp \psi$ *or* $\phi' \perp \tau'$.

$\langle t, t \rangle$ *is said to be externally inappropriate in* L_1 *(wrt* Ψ*) if there is some channel* L_2 *such that* $\langle t, t \rangle$ *is externally inappropriate in* L_1 *and* L_2 *(wrt* Ψ*).*

(The type-level conditions specified in the above definition are illustrated in Figure 5.2.) The condition that the tuple is not already internally inappropriate in one of the links ensures that normality is indeed *maximal*—i.e. if $\langle t, t \rangle$ is internally inappropriate in L_1, then it will not support the corresponding inference and so the inference supported by L_2 can take place without fear of a contradictory conclusion being reached. Similarly—also for reasons of maximising normality—the only occasions in which external inappropriateness needs to be imposed upon two links is if the antecedents of the links are actually satisfied (or if the antecedent of one is satisfied and the antecedent of the other is entailed by the information thereby inferred). For example, there is no need for *tweety* $\mapsto$ *tweety* to be externally inappropriate in links supporting

$peacocks \rightarrow bright.feathers$

$peacocks \rightarrow lays.eggs$

if *tweety* is not a peacock.

Having defined the above concept, the concept of *inappropriateness* is extended as

[18] I will often abbreviate this by saying "$\langle t, t \rangle$ is externally inappropriate in L_1, L_2".

$$\begin{tikzcd}[column sep=small] L'_1 : \phi'\rightarrow\psi' \arrow[dr, dash, "f_1"] & & L'_2 : \tau'\rightarrow\sigma' \arrow[dr, dash, "f_2"] & \\ & L_1 : \phi\rightarrow\psi & & L_2 : \tau\rightarrow\sigma \end{tikzcd}$$

$$\phi' \perp \sigma \text{ or } \tau' \perp \psi \text{ or } \phi' \perp \tau'$$

Figure 5.2: Conditions resulting in external inappropriateness

follows: $\langle t, t\rangle$ is *inappropriate* in $L^{\wedge}$ wrt Ψ if either $\langle t, t\rangle$ is internally inappropriate in $L^{\wedge}$ wrt Ψ or $\langle t, t\rangle$ is externally inappropriate in $L^{\wedge}$ wrt Ψ. Given this redefinition of inappropriateness, the definition of maximal normality is unchanged—it simply involves this extended version of inappropriateness.

Finally, as was done with the Reliability Constraint in Section 3.5.5, the definition of external inappropriateness is generalised to cover the case whereby more than two links are mutually conflicting.

Definition Revised external inappropriateness: *Let* $L_1 : A \rightrightarrows A$, *...,* $L_n : A \rightrightarrows A$ *be information links such that* $\langle \phi_1, \psi_1\rangle \in L_1^{\wedge}$, *...,* $\langle \phi_n, \psi_n\rangle \in L_n^{\wedge}$. *Further let* Ψ *be a collection of propositions formed from the types and tokens of* A *and let* $t \in tok(A)$ *be such that* $(t : \phi_1)$ *,...,* $(t : \phi_k)$ *all hold in* Ψ *and* $(\phi_{k+1} \wedge \ldots \wedge \phi_n) \leq_A (\psi_1 \wedge \ldots \wedge \psi_k)$. *The pair* $\langle t, t\rangle$ *is said to be* externally inappropriate to the signalling relations $L_1^{\wedge}$,..., $L_n^{\wedge}$ wrt Ψ *if:*

1. *$\langle t, t\rangle$ is internally inappropriate in neither L_1 nor ... nor L_n wrt Ψ;*
2. *there exist links L'_1, ..., L'_n such that $L_i \preceq_{f_i} L'_i$ for all i, $f_i(\langle \phi_i, \psi_i\rangle) = \langle \phi'_i, \psi'_i\rangle$ for all i, and $\bigwedge_i \xi_i$, where ξ_i is either ϕ'_i or ψ_i, is an inconsistent type.*

The intuitions behind this definition are identical to those for the simpler definition above; the revised definition simply acknowledges that a given set of types may be mutually conflicting even when no pair from this set is pairwise conflicting.

5.3.3 Channel Theoretic Consequence Relations

A consequence relation is a collection of pairs of sentences. By constraining the set of initial information (on which default reasoning is performed) to be a singleton set, $\mid\!\!\leadsto_{\mathcal{S}}$—where $\mathcal{S}$ is a collection of channels—can similarly be seen as a relation between propositions, defined as follows:[19]

> $p \mid\!\!\leadsto_{\mathcal{S}} q$ iff $q \in Consq_{\mathcal{D}}(\{p\})$, where $\mathcal{D}$ is the default inference system based on $\mathcal{S}$ and $\{p\}$.

Whether one would want to call $\mid\!\!\leadsto_{\mathcal{S}}$ a consequence relation is debatable since it is not a relation between sentences; however, $\mid\!\!\leadsto_{\mathcal{S}}$ is a relation that models inference in the channel-theoretic framework, and this is the important property. For convenience, I will call such a relation a *default inference relation.*

The *KLM* conditions on consequence relations can be modelled in the channel-theoretic setting by defining analogous conditions that a "reasonable" default inference relation should satisfy. By interpreting propositional connectives via type operations,[20] the *KLM* principles can be given a channel-theoretic interpretation in a straightforward manner. This is illustrated by the modification of *LLE* below. The other principles are similarly modified to obtain a collection of constraints applicable to default inference relations. As usual, standard material implication is modelled by inference via the logical link. The fact that $(a : \phi) \mid\!\!\leadsto_{\mathcal{S}} (a : \psi)$ is supported by the logical link L—i.e. $\langle a, a\rangle \in L^{\vee}$ and $\langle \phi, \psi\rangle \in L^{\wedge}$—is written as follows: $(a : \phi) \models\!\!\!\Rightarrow_L (a : \psi)$. I will drop the subscripts from $\mid\!\!\leadsto$ and $\models\!\!\!\Rightarrow$ when this causes no confusion.

Revised LLE: $$\frac{(a:\alpha) \models\!\!\!\Rightarrow (a:\beta) \qquad (a:\beta) \models\!\!\!\Rightarrow (a:\alpha) \qquad (a:\alpha) \mid\!\!\leadsto (a:\gamma)}{(a:\beta) \mid\!\!\leadsto (a:\gamma)}$$

This condition is interpreted in the following way: if a given relation $\mid\!\!\leadsto$ satisfies the conditions above the horizontal line, then it must also satisfy the condition below the line. The other principles of preferential and rational consequence relations are similarly modified. The central result of this section is that the channel-theoretic model of default inference satisfies the conditions corresponding to the *KLM* principles associated with

[19]Since I have assumed Ψ is always a singleton set of the form $\{p\}$, I will tend to write '$p \mid\!\!\leadsto_S q$' rather than '$\{p\} \mid\!\!\leadsto_S q$'.

[20]This can be done since the signalling relations are reflexive—e.g. $(a : \phi) \wedge (a : \psi)$ holds iff $(a : \phi \wedge \psi)$ holds, where the first instance of '$\wedge$' denotes proposition-conjunction and the second is type-conjunction.

preferential consequence.

Theorem 5.3.2 *For any collection $\mathcal{S}$ of channels, $\mid\!\sim_{\mathcal{S}}$ satisfies all the conditions corresponding to the principles required of preferential consequence relations.*

Proof See Appendix D. $\square$

This result shows that the progression through the conditional channel operations, the specific definition of a default hierarchy, and the imposition of a normality condition leads to a model of default reasoning satisfying important principles regarding which inferences are supported. Unfortunately, the extension to rational relations is not straightforward—since the principle of Rational Monotonicity involves the assertion that a pair of sentences is *not* in a consequence relation, then a proof that Rational Monotonicity is supported is more problematic.[21] However, the intuition behind the definition of a kernel default hierarchy (Section 4.3.3)—whereby a type α is a background condition of a constraint with antecedent β only if there is a constraint linking β to $\neg\alpha$—should be sufficient to ensure that Rational Monotonicity is supported. While I conjecture that the channel-theoretic system of default reasoning models rational consequence relations, this problem unfortunately remains open at this stage.

5.4 Towards a Methodology for Situated Default Reasoning

In this section, I propose a methodology for default reasoning that represents a subtle variation to the usual approach. The methodology is based on the hierarchical channel theoretic framework I used throughout the thesis. The work described in this section is still embryonic at this stage, but seems to offer an interesting avenue for future investigation.

Default reasoning is usually thought of as reasoning with incomplete information. By contrast, the methodology I outline here can be thought of as reasoning with incomplete, or approximate, constraints. I have described before how inference via a channel or link is inference at some "level of approximation"—there may be a channel higher up

[21] Actually, Lehmann (1989) shows that if the logical language is restricted so that negations of conditional assertions are disallowed, then preferential consequence is equivalent to rational consequence.

the hierarchy which is potentially "more reliable" than the chosen one, but the agent makes a particular choice (with respect to which channel to use for a particular inference) based on the chosen level being considered appropriate for the task at hand.[22] This view leads to a methodology for default reasoning which involves a subtle shift from the more standard view but seems to offer the potential for a more natural interpretation of commonsense reasoning. At this point, I only sketch the ideas for such a methodology and outline how it relates to the formal model of default reasoning described in this chapter. Section 5.4.3 provides a bit more detail to the basic ideas within the framework of the *qualification problem* in AI planning (Georgeff 1987b), which is traditionally handled using default reasoning (e.g., see Shoham (1988)). A more comprehensive investigation of the methodology is earmarked in Section 6.2 for future research.

5.4.1 Outline of the Methodology

The methodology relies on a hierarchy of channels supporting various constraints. As usual, for channels $\mathcal{C}$, $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq \mathcal{C}'$, $\mathcal{C}'$ is more reliable than $\mathcal{C}$ in that a connection c that is an exception in $\mathcal{C}$ need not be an exception in $\mathcal{C}'$, yet every connection c' that is a token in both channels and is an exception in $\mathcal{C}$ must also be an exception in $\mathcal{C}'$. In the particular model of default reasoning I have in mind, each level in the hierarchy corresponds to a particular "refinement" of a regularity—moving up a level in the hierarchy involves adding conditions to the antecedent and/or consequent of constraints, thereby resulting in a more informative version of the associated regularity.[23]

Given certain information which is to be used as the premise of some inference, the choice arises as to which level in the hierarchy is the most appropriate for the inference. This choice involves a trade-off. Using a channel higher up in the hierarchy results in a lower likelihood of error, since the connection in question is less likely to be an exception. However, choosing a channel lower in the hierarchy results in simpler reasoning, since the classifications linked by the subchannel will have a less discriminating set of types

[22]The way such a decision is made is, of course, a non-trivial problem! The main issue to be considered involves the trade-off between a simpler constraint (i.e. one in a channel further down the hierarchy) and thereby simpler inference, versus a more informative constraint (i.e. one in a channel higher up the hierarchy) and thereby increased reliability in the form of fewer exceptions. This trade-off is discussed in the following but is by no means resolved.

[23]Note that while I talk of the constraint from higher up the hierarchy as being a "more informed version" of the other, I should stress, as I have throughout, that both constraints correspond to *bona fide* regularities in their own right.

than those linked by the super-channel.

Consider the following simple example. The generic "Birds fly" describes a channel $\mathcal{C}$ containing a constraint of the form

$$bird \rightarrow flies.$$

The associated hierarchy also contains a channel $\mathcal{C}'$, such that $\mathcal{C} \sqsubseteq \mathcal{C}'$, supporting a constraint of the form

$$bird \wedge \neg penguin \rightarrow flies.$$

Given an individual *tweety* and information $(tweety : bird)$, one must decide whether to project this information using the link $Link(\mathcal{C})$ or the link $Link(\mathcal{C}')$. This effectively involves choosing whether the connection $tweety \mapsto tweety$ should be in $tok(\mathcal{C})$ or only in $tok(\mathcal{C}')$. If $(tweety \mapsto tweety) \in tok(\mathcal{C})$ and $Link(\mathcal{C})$ is used for the inference, then there is a greater chance of the connection being an exception and the projection leading to an erroneous conclusion. However, it would also allow the inference to $(tweety : flies)$ without having to consider whether or not *tweety* is a penguin.

In the spirit of default reasoning, the choice of appropriate channel is made by considering what other information is known about *tweety*. If $(tweety : penguin)$ holds, then clearly $\mathcal{C}'$ is the more appropriate channel. This leads to the following notion of *adequacy*.[24]

Definition *Given a collection Ψ of propositions, a channel $\mathcal{C} : A \Rightarrow A$ is* adequate with respect to Ψ *if for each $(t : \psi) \in \Psi$, $t \in tok(A)$ and $\psi \in typ(A)$.*

A channel $\mathcal{C}$ is adequate with respect to a collection Ψ of propositions if the associated classification A contains all the types and tokens that are used in the propositions in Ψ—effectively, the "language" of A is rich enough to represent all the propositions of Ψ. To ensure that the simplest channel is chosen, the appropriate channel for a given inference is taken to be a minimal adequate one for the propositions in question.

By requiring the appropriate channel to be adequate, the potential problem with penguins is ruled out—if $(tweety : penguin)$ is known to hold, then the channel $\mathcal{C}$ containing $bird \rightarrow flies$ is not deemed appropriate. The model of default reasoning this leads to is one of choosing the context (i.e. channel) that is most appropriate to the problem

[24] I will continue to adopt the simplification that all channels are reflexive. The following definition would have to be revised if this simplification was lifted.

at hand and reasoning monotonically within that context. This seems to lead to a more natural approach to commonsense reasoning, which is demonstrated by application to the qualification problem in Section 5.4.3.

Of course, the above outline of a methodology for defeasible reasoning is just that—an informal outline. It could be that any gains—conceptual or computational—are simply balanced by the burden of determining the appropriate (i.e. minimal adequate) channel on which to base the inference. In practice, it seems that we make choices as to which level of discrimination is most appropriate in ways that may be impossible to capture formally. However, at this point I believe that the approach has the potential to lead to a procedure for reasoning with incomplete information that is more natural and possibly more computationally efficient than those based on more standard default logics.

5.4.2 Relation to Maximal Normality Model

An important property of the methodology for defeasible reasoning I have outlined is the way in which it relates to the concept of maximal normality of Section 5.2.2.[25] In fact, the methodology suggested itself from the maximal normality definition rather than independently. As above, the discussion here is an outline of an idea which requires further investigation.

The maximal normality condition requires each connection in each channel of a hierarchy to be (internally) appropriate. A connection $t \mapsto t$ is inappropriate, with respect to some given information $(t : \phi)$, if $\phi \perp \phi'$ for some channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq \mathcal{C}'$ and ϕ' is the antecedent of a constraint in $\mathcal{C}'$.[26] Informally speaking, channels higher up the subchannel ordering have to be checked before determining whether a connection c is to be included in the token-set of a given channel or not. Of course, if c *is* included in $tok(\mathcal{C})$, then it will support an inference via $Link(\mathcal{C})$ (assuming that the source is classified by the antecedent of the constraint) regardless of whether it supports an inference via $Link(\mathcal{C}')$ for any super-channel $\mathcal{C}'$ of $\mathcal{C}$. This corresponds to the case where the "simpler" channel is used to perform a defeasible inference in the methodology above.

[25]For simplicity of this discussion, I will only consider the initial definition of maximal normality—i.e. that involving internal inappropriateness only (but see the comment in Appendix B.2).

[26]The inappropriateness condition also involves the consequents of constraints but, again for simplicity, I restrict discussion to antecedents here.

Consider now the case of $t \mapsto t$ being inappropriate in $\mathcal{C}$. It must be the case that $(t : \phi)$ holds, as does the property described in the previous paragraph. In this case, $t \mapsto t$ is not inserted as a token of $\mathcal{C}$ since it is inappropriate. However, $t \mapsto t$ is, by definition, *not* inappropriate in $\mathcal{C}'$ since $(t : \phi')$ cannot hold, where ϕ' is the antecedent of the offending constraint in $\mathcal{C}'$. This corresponds to $\mathcal{C}'$ being the *adequate* level in the hierarchy in which to reason about the information at hand. An example helps to make this more explicit.

Consider the usual *penguin* example, with $\mathcal{C}$ containing the constraint

$$bird \rightarrow flies$$

and $\mathcal{C}'$ containing the constraint

$$bird \wedge \neg penguin \rightarrow flies$$

and suppose that $(tweety : bird)$ and $(tweety : penguin)$ hold. In the methodology outlined above, $\mathcal{C}$ is inadequate with respect to this information since it does not deal with information related to the type *penguin*.[27] Similarly, the connection $tweety \mapsto tweety$ is inappropriate in $\mathcal{C}$ because information known about *tweety* conflicts with the antecedent of the constraint in $\mathcal{C}'$. Hence, *tweety* is a pseudo-signal in $\mathcal{C}$—i.e. the connection $tweety \mapsto tweety$ is, by definition, not contained in $tok(\mathcal{C})$.

The procedure does not exactly match the abstract definition. For example, if the information $(tweety : \neg penguin)$ was known, then the connection $tweety \mapsto tweety$ would not be inappropriate in $\mathcal{C}$ but $\mathcal{C}$ would be inadequate with respect to this information. In this case it would of course be safe, and appropriate, to use $\mathcal{C}$ as the channel on which to base the inference regarding *tweety*'s flying abilities. This suggests that the definition of adequacy should take into account whether the given information is "relevant"—i.e. whether having that information is good reason for choosing a higher level in the hierarchy as the most appropriate one. However, this would complicate the adequacy definition somewhat.

[27] I am tacitly assuming that the type *penguin* is not contained in the type-set associated with the classification A linked by $\mathcal{C}$. This assumption is formalised in Section 5.4.3.

5.4.3 An Application of Default Reasoning: The Qualification Problem

In this section, I sketch out in more detail the methodology for default reasoning outlined above by considering the *qualification problem* from the AI planning literature (e.g. (Georgeff 1987b; Ginsberg and Smith 1988)), whereby not all necessary preconditions of an action can be easily specified. The solution to this problem is usually cast in terms of some version of default logic. I originally formulated the approach below independently of any other result described in this thesis—in fact, it was the first problem for which I used the concept of a hierarchy of channels—but have since come to see it as a particular instance of the more general methodology described earlier.[28] It is hoped that further work on this problem will strengthen the claim that channel theory, and particularly the hierarchy framework described in this thesis, provides a powerful unifying framework for reasoning about conditionals, defaults and plans.

In the simple model of events I employ here, an event is simply a connection between two situations. This is consistent with the standard AI formal representation, based on McCarthy and Hayes' (1969) *situation calculus,* whereby an event involves a transition between world states, which are instantaneous snapshots of the world. Representing and reasoning about events is an important area of AI research and much work has been dedicated to providing adequate representations of events, their preconditions and effects, and related issues such as causation and synchronisation, within the situation calculus framework (e.g. see (Georgeff 1987b; Allen et al. 1990).[29] The representation of events I define below ignores most of these issues, focussing solely on the qualification problem.

Outline of a Simple Model of Events

In the standard AI model of events, each event has associated with it a number of *preconditions* and *postconditions* (or *effects*). If $pre(e)$ denotes the preconditions of an event e and $post(e)$ the postconditions, then the general interpretation given them is the following:

[28] A related approach to the qualification problem, within a classical logic framework extended by a notion of *context,* has recently been proposed by Bouquet and Giunchiglia (1994).

[29] A particular problem that has achieved notoriety is the *frame problem,* but I will not be tackling that here.

> If e occurs in a state in which all of $pre(e)$ hold, then the resulting state immediately following the performance of e will satisfy all of $post(e)$.

The situation calculus is a framework devised for formalising this constraint. It is basically a sorted first-order logic with three sorts: *fluents* (terms that denote conditions on world states), *events* and *situations.* For example, a formula representing the above constraint could be written as follows:
$\forall e \forall s((\forall p(pre(e,p) \supset holds(p,s)) \wedge occurs(e,s)) \supset \forall p(post(e,p) \supset holds(p, succ(s))))$
where $pre(e,p)$ denotes that p is a precondition of e, $post(e,p)$ denotes that p is a postcondition of e, $occurs(e,s)$ denotes that e occurs in s, $holds(p,s)$ denotes that p holds in s, and $succ(s)$ denotes the successor state of s. (Lifschitz (1987) outlines a general methodology for representing constraints of the above form in the situation calculus.)

The channel-theoretic model of events I use here is a very simple one.[30] An event is a connection in a particular channel. Different channels correspond to different *event-types,* with a given connection corresponding to a particular instance of that event-type. For example, there may be a *Move* channel such that every connection in $tok(Move)$ corresponds to a "move" event. I will assume that all sources and targets of event-connections are situations, ensuring a correspondence with the view described above whereby events involve transitions between situations. A channel-type is the constraint between the corresponding event-type's pre- and post-conditions. I will call such a type an *event-description.* For example, if the *MoveA* channel models the event-type of moving object A from the door to the window, then the corresponding constraint could be of the form

$$(at(A, door) \wedge \neg heavy(A)) \rightarrow at(A, window).$$

Of course, to represent more general event-types I would have to make use of parameters. However, for simplicity, I will consider only non-parametric objects.

A sequence of events is modelled via serial composition of channels—the complex event-type is the composition of the two channels while the complex event is the connection that results from composing the connections corresponding to the original events.[31]

[30] There are obvious parallels between this simple model of events described and Barwise's (1993) channel-theoretic model of Hoare logic.

[31] For simplicity, I will assume that standard composition, rather than conditional composition, is appropriate. However, if it is assumed that the source and target of an event-connection are always distinct (this would be the case if every event had some effect on a situation), then these two forms of

It should also be possible to use some form of parallel composition to model concurrent events. However, this is more problematic than sequential events and I do not pursue it here. Causation (i.e. the occurrence of one event causing the occurrence of another) and ramification (i.e. side-effects of an event) can be modelled using other channels. The fact that the target situation of an event is the temporal successor[32] of the source situation can also be modelled using a channel $\mathcal{T}$, whose collection of tokens ranges over all events, and which contains a constraint that indicates temporal succession between source and target. Enforcing *frame axioms*—i.e. axioms that state that nothing changes during the performance of an event except for those conditions contained in the postcondition of the event—is more problematic. One approach is to have an extra constraint in each event-channel, indicating that (for each connection) the target supports all the information supported by the source except for that which contradicts the postconditions of the related event-type (given by the consequent of the corresponding constraint). It is not clear exactly how this would be stated and it would certainly require the use of parameters or n-ary types.[33] In any case, I will not pursue these issues at all—my concern here is to outline a solution to the qualification problem using the methodology for defeasible reasoning described earlier.

The Qualification Problem

The qualification problem concerns the issue that it is impossible to completely list all the conditions that need to hold for an action or event to be performed successfully (i.e. to have the desired effects). For example, if I turn the ignition key in my car, it usually starts. However, it may not do so on one particular morning for any of a number of reasons: for example, the battery may be flat or disconnected, or the starter-motor may be burnt out, or the engine may be flooded. This problem is, of course, well known to philosophers and accounts for the nonmonotonicity of conditional sentences (e.g. "If I strike the match then it will light" does not hold if the match is wet, or the matchbox is wet, or the match has already been used, etc.). In Chapter 3, I defined a framework

---

composition are equivalent anyway.

[32]Requiring an event to involve discrete temporal succession is, of course, a simplification. However, this suffices for the purposes of the current discussion.

[33]Georgeff et al.'s (1993) partial logical framework for reasoning about action allows for different events to occur simultaneously, with some of these events being "unknown" to the agent. In particular, this means that the world may change in unexpected ways while an action is being performed and, on completion of that action, no conditions can be assumed to hold except for the postconditions of the action just performed. Dropping frame axioms achieves a similar effect.

for reasoning about conditionals in which the background assumptions underlying a conditional are represented via a hierarchy of channels. This is the obvious approach to take in addressing the qualification problem within the channel theoretic framework.[34]

The usual approach to qualification involves the use of a default logic in some way. For example, consider the approach taken by Shoham (1988). It is not necessary to describe Shoham's particular nonmonotonic logic of *chronological ignorance*, but simply to outline his approach. Shoham formalises actions using a temporal logic augmented with a modal operator '$\Box$', where $\Box\phi$ represents "ϕ is known". In Shoham's framework, the rule for striking a match is represented as follows:[35]

$$\begin{aligned}
&\Box holds(holding.match, s) \ \wedge \ \Box occurs(strike.match, s) \wedge \\
&\qquad \neg\Box holds(wet.match, s) \wedge \\
&\qquad \neg\Box holds(wet.matchbox, s) \wedge \\
&\qquad \neg\Box holds(used.match, s) \ \wedge \ \ldots \\
&\quad \supset \Box holds(lit.match, succ(s)).
\end{aligned}$$

The nonmonotonic aspect of Shoham's logic ensures that $\neg\Box\phi$ is a consequence of a theory whenever $\Box\phi$ is not.

There are a couple of problems with Shoham's solution (and the default logic approach in general). To reason about the effect of an action, an agent still needs to work through a checklist of possible contingencies—the background assumptions still need to be explicitly listed and the agent cannot reason about whether an action will be successful without considering each of these assumptions. The default logic ensures that it is enough to not have information that the assumptions are violated, rather than requiring information that the assumptions are not violated. However, this still does not seem reasonable. When I turn the key to start my car, I do not wonder whether the battery is flat, and—not having information to the contrary—assume that it is not. I turn the key and expect (hope!) the car to start.[36] In Shoham's formalisation, the regularity

[34] After all, an event channel can be seen as supporting a particular sort of conditional sentence. For example, the *Move* channel of the previous section supports conditionals of the form:
"If in situation s A is at the door and A is not heavy, then (after performing the *move* event) in situation s' A is at the window."

[35] I have modified the language slightly so that it is more consistent with the situation calculus treatment. However, this is not significant. (For some reason, Shoham treats events as if they were propositions.)

[36] In some situations, for example, cold mornings, the expectation may be lowered and I may decide to explicitly consider some background assumptions.

representing the behaviour of the *strike.match* or *start.car* events explicitly incorporates the background assumptions. However, a more realistic approach is to use a regularity that has such background assumptions as implicit, unstated conditions—i.e. if I strike the match, then it will light. A regularity such as this is fallible, of course—it may be the case that one of the background assumptions is violated, even though I do not know it—but it is reliable enough to be dependable in normal situations. Of course, this is exactly what the channel theoretic solution provides: a means of modelling reliable yet fallible event behaviours.

In the next section, I define a framework for using fallible regularities to reason about events based on the channel hierarchy framework that I have used throughout the thesis. I then outline a methodology for reasoning about the expected effects of events, including the possibility of an event failing to have its expected effect.

A Channel Theoretic Model

An *event-channel* is a channel $\mathcal{E} : A \Rightarrow A$ whose tokens are connections between two situations s and s' such that s' is the temporal successor of s, and whose constraints link preconditions of the associated event-type to postconditions of that event-type. A *model* of an event-type is a collection of event-channels arranged in a channel-hierarchy. Moving up the hierarchy corresponds to "refining" the description of the event-type, i.e. making implicit background conditions explicit. For example, one channel for starting a car may contain a constraint of the form

$$key.turned \rightarrow car.started$$

whereas a channel higher up the ordering may instead contain the constraint

$$key.turned \wedge battery.ok \rightarrow car.started.$$

This is, of course, the way in which I have used channel hierarchies to encode implicit background assumptions throughout the thesis. The following definitions serve to capture the intuition that each level in a hierarchy captures a particular approximation (i.e. using a restricted set of types) of a regularity,[37] and that moving up the hierarchy allows

[37] Strictly speaking, I don't like to think of any of the channels in a hierarchy supporting an "approximation" of some regularity whose "correct" form is contained higher up the hierarchy. Each channel supports a regularity—the fact that a superchannel supports a more robust version of that regularity does not detract from the original one being a regularity in itself. I have tried to stress this point throughout the thesis. That said, the particular language I have used is (while potentially misleading)

the set of types to be extended and therefore results in a more accurate approximation of the regularity.

The first definition ensures that moving up a level makes a regularity more informative.

Definition *A channel* $\mathcal{C} : A \Rightarrow B$ *is* redundant *in a hierarchy if for all constraints* $\gamma \in typ(\mathcal{C})$ *and some distinct subchannel* $\mathcal{C}' : A' \Rightarrow B'$ *such that* $\mathcal{C}' \sqsubset \mathcal{C}$*, it is the case that* $ante(\gamma) \in typ(A')$ *and* $succ(\gamma) \in typ(B')$*.*

Intuitively, a channel $\mathcal{C}$ is redundant if there is a subchannel $\mathcal{C}'$ in the hierarchy such that the types available to $\mathcal{C}'$ are sufficient to represent the constraints in $\mathcal{C}$. This ensures that any new level in the hierarchy "adds" information, in that at least some regularity has had a new type—one that wasn't available at any level below that channel—added to its antecedent or consequent. This is intended to capture the notion of a channel being a representation of a regularity at a particular "level of approximation".

The second definition ensures that moving up a level "conservatively extends" the regularities—i.e. each regularity at a given level is identical to the corresponding regularity at any lower level if restricted to the types available at the lower level.

Definition *A channel* $\mathcal{C} : A \Rightarrow B$ *is* conservative *in a hierarchy if for each constraint* $\gamma \in typ(\mathcal{C})$ *and each distinct subchannel* $\mathcal{C}' : A' \Rightarrow B'$ *such that* $\mathcal{C}' \sqsubset_f \mathcal{C}$*, there is a constraint* $\gamma' \in typ(\mathcal{C}')$ *such that* $f(\gamma') = \gamma$ *and*

1. $ante(\gamma') = ante(\gamma)$*, or* $ante(\gamma) = (ante(\gamma') \wedge \phi)$ *for some* $\phi \in (typ(A) \setminus typ(A;))$*;*
2. $succ(\gamma') = succ(\gamma)$*, or* $succ(\gamma) = (succ(\gamma') \vee \psi)$ *for some* $\psi \in (typ(A) \setminus typ(A'))$*.*

The final restriction required is that the same set of events (i.e. connections between situations) is the subject of all the channels in the hierarchy—i.e. each channel in a model of an event-type ranges over the same collection of events.

Definition *An* event-model $\mathcal{M}$ *for an event-type is a hierarchy of channels (event-channels)* $\{\mathcal{E}_1, \mathcal{E}_2, ...\}$*, where* $\mathcal{E}_i : A_i \Rightarrow B_i$*, such that*

1. $tok(\mathcal{E}_i) = tok(\mathcal{E}_j)$ *for all* i, j*;*
2. *no* $\mathcal{E}_i$ *is redundant; and*
3. *every* $\mathcal{E}_i$ *is conservative.*

---

rather convenient.

Given the intuitive interpretation of an event-model, it seems reasonable to assume that the hierarchy satisfies certain other conditions. For example, every pair of channels should have a common subchannel—this prevents one event-model from containing event-channels corresponding to different event-types. Further, one may want to define meet and join operations on channels and require the hierarchy to be a complete lattice. However, I will not impose any such requirements.

Given two event-models, serial composition leads to a new event-model (with the original channel hierarchies effectively merged). Thus, given a number of event-models, a full hierarchy of event-channels can be obtained via repeated composition. I will not describe this process in detail here but simply note that sequences of events can be modelled in exactly the same manner as "primitive" events.[38]

Reasoning about the effects of events is, of course, a matter of prediction. In the channel-theoretic framework, this is modelled by projection along information links, in the manner described in Sections 2.3.6 and 5.2.1. As described above, choosing a link (with which to reason with) from a link hierarchy requires determining the appropriate level in the hierarchy. This was informally captured using the concept of *adequacy.* Inference then proceeds via that link. Once again, I use the following notation: given a link L and proposition $(s : \phi)$, I write $(s : \phi) \mathrel{|\!\!\rightsquigarrow}_L (s' : \psi)$ iff there is a link L (in some given collection) such that $\langle s, s' \rangle \in L^\vee$ and $\langle \phi, \psi \rangle \in L^\wedge$.

Definition *Let E be an event-type with associated event-model $\mathcal{M}$ and $(s : \phi)$ a proposition. The* predicted effects of an event of type E in s *is the collection of propositions Ψ defined as follows:*

$$\Psi = \{(s : \psi) \mid (s : \phi) \mathrel{|\!\!\rightsquigarrow}_{Link(\mathcal{E})} (s : \psi),\ \textit{where } \mathcal{E} \in \mathcal{M} \textit{ is the minimal adequate event-channel for } \phi\}.$$

The appropriate level in a hierarchy is determined by the notion of adequacy. However, choosing the premise proposition—i.e. determining which facts about the initial situation should be checked before being able to predict the outcome of an event with any confidence—is an extremely problematic issue and well outside the scope of the model described here. For example, when I go out to my car in the mornings, I tend not to check whether the battery is connected (i.e. I don't check whether $(s : \textit{battery.connected})$ holds). However, if I came across an old rusted model in an out-of-the-way barn, I may

[38] As well as the serial composition of event-channels, a full channel hierarchy would contain channels that modelled causation and ramification of events, as outlined earlier.

well want to check the battery connection (as well as other things!) before trying to turn the engine over.

Failure and Recovery

When using an approximation of a regularity to reason (defeasibly) about the effects of events, we open up the possibility of making erroneous predictions. For example, one morning when I go out and turn the key in the ignition of my car, I may find that the car indeed fails to start. This could be for any of a number of reasons, each of which is a background assumption to the channel $\mathcal{C}$ supporting

$$key.turned \rightarrow car.started.$$

Such an event is modelled as an exception in $\mathcal{C}$.

For example, consider the situation s where I go out to my car one sunny morning. As usual, I expect my car to start on turning the key in the ignition. I therefore use the link $Link(\mathcal{C})$, where $\mathcal{C}$ supports the above regularity, to reason that my car will be running in the (temporally successor) situation s', where $s \mapsto s'$ is a connection in $\mathcal{C}$. However, for some reason, I find that my car has not started in s'. This must mean that $s \mapsto s'$ is an exception in $\mathcal{C}$, of course—if it were not, then (by the Principle of Harmony) it would be the case that (s : *key.turned*) and (s' : *car.started*) both held. However, as shown in Section 3.3.2, an exception to $\mathcal{C}$ need not be an exception to some channel higher up the hierarchy. For example, suppose that the reason the car didn't start was because the battery was disconnected (my wife playing a practical joke!), and suppose that this condition is encoded as a background assumption.[39] Then there is a channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, where, say,

$$f(key.turned \rightarrow car.started) = \quad key.turned \wedge battery.connected \rightarrow car.started.$$

In this channel, $s \mapsto s'$ is *not* an exception since (s : *key.turned* $\wedge$ *battery.connected*) does not hold in the classification A' containing this type. This leads to the following definition, modelling the error-recovery process.

[39] As discussed in Chapter 3, a channel hierarchy can be seen as a parameter in the representation of an agent, encoding the implicit background assumptions held by a particular agent—i.e. different agents (or the same agent in a different state) may have different background assumptions regarding the "same" regularity.

Definition *Let $\mathcal{M}$ be an event-model, $\mathcal{E} : A \Rightarrow B$ an event-channel from $\mathcal{M}$, $e \in tok(\mathcal{E})$ an event-connection and $\gamma \in typ(\mathcal{E})$ an event-description such that e is an exception to γ in $\mathcal{E}$. A channel $\mathcal{E}' : A' \Rightarrow B'$ is a* solution to $\langle \mathcal{E}, (e : \gamma) \rangle$ in $\mathcal{M}$ *if*

1. *$\mathcal{E} \sqsubseteq_f \mathcal{E}'$ in $\mathcal{M}$; and*
2. *e is not an exception to $f(\gamma)$ in $\mathcal{E}'$.*

$\mathcal{E}'$ is a preferred solution if there is no other solution $\mathcal{E}''$ such that $\mathcal{E}'' \sqsubseteq \mathcal{E}'$ in $\mathcal{M}$.

The above definition models the following method of dealing with error. If an agent makes a prediction, using $\mathcal{E}$, that an event will have a particular effect and then finds that the successor situation does not satisfy the predicted condition, then the minimal channel for which that event is not an exception contains the "reason" why the event failed. For example, in the case of starting my car, the channel whose event-description contains *battery.connected* in the antecedent "describes" why turning the key failed to start my car—i.e. the battery was disconnected. (I don't mean this to be a claim regarding how I go about determining why my car failed to start, of course. Describing such a procedure, which would rely on all sorts of heuristics about my car and the noises it makes, is well outside the scope of this model.)

It could be the case that there is *no* solution to a given exception. This corresponds to the case where the problem that caused the corresponding event to fail to have its predicted effects is not encoded as part of the event-model—i.e. that particular background assumption is not represented in the channel hierarchy. For instance, if I know very little about cars, then I may not realise that the distributor not being faulty is a background assumption behind my car starting. Alternatively, a condition that I may never have even considered, simply because it was ludicrous, may be behind the car's failure—for example, my wife may have played another of her practical jokes, inserting a potato into the exhaust-pipe. Such cases—where the implicit conditions are insufficient to explain the exception—correspond to the opportunity for the agent to expand its knowledge by adding to the hierarchy.

Since I am talking about reasoning and inference, I should really be using links rather than channels. This involves using the link hierarchy constructed from an event-model hierarchy, as defined in Section 5.2.1, and modifying the above definition in the obvious way, replacing channels by links. The following results relate error-recovery under links to error-recovery under channels.

Proposition 5.4.1 *If the channel $\mathcal{C}$ is a solution to $\langle \mathcal{E}, (e : \gamma)\rangle$ then the link $Link(\mathcal{C})$ is a solution to $\langle Link(\mathcal{E}), (e : \gamma)\rangle$.*

Proof Obvious, since any exception to $Link(\mathcal{C})$ is also an exception to $\mathcal{C}$, and by Proposition 5.2.1, $Link(\mathcal{C}) \preceq Link(\mathcal{C}')$. $\square$

Proposition 5.4.2 *For each channel $\mathcal{C}$ that is a preferred solution to $\langle \mathcal{E}, (e : \gamma)\rangle$, there is a link L that is a preferred solution to $\langle Link(\mathcal{E}), (e : \gamma)\rangle$ such that $L \preceq Link(\mathcal{C})$.*

Proof By the previous result, $Link(C)$ is a solution, so by definition there must be a solution L satisfying the required property. $\square$

The above result seems to suggest that there may be a "better" solution from amongst the links than from amongst the channels. This is because an event may fail to have its predicted effect, but the predicted condition still holds due to some other influence. For example, Fred dies when I pull the trigger of the gun, even though the gun isn't loaded, because he has a cardiac arrest (Fred clutching his chest hides the fact from me that he isn't bleeding). However, in the absence of such coincidences—i.e. *weak exceptions* (defined in Section 2.3.4)—solutions amongst the channels and links coincide.

Proposition 5.4.3 *Let $\mathcal{C}$ be a channel containing no weak-exceptions. Then $\mathcal{C}$ is a solution to a given exception if and only if $Link(\mathcal{C})$ is.*

Proof Exceptions in $\mathcal{C}$ exactly correspond to exceptions in $Link(\mathcal{C})$ in the absence of weak exceptions. $\square$

5.5 Discussion

The channel theoretic analysis of generics of the previous chapter involved a generic being simply true or simply false, regardless of the properties of the associated individuals. This necessitated an extra assumption—one of maximal normality—to be imposed on the signalling relations connecting tokens before the channel theoretic framework for generics could be used for default reasoning about the properties of individuals. What is of greater interest, however, is the way in which the specific maximal normality condition

I defined, particularly the way in which it is constructed with respect to a channel hierarchy, relates to a model of situated reasoning in which an "appropriate" regularity—appropriate with respect to known information—is picked from a hierarchy and is used as the basis of inference. One of the tenets of situated reasoning is that a regularity can be used without necessarily explicating all the background assumptions required to be in place for the regularity to be reliable. As such, the methodology for reasoning outlined in Section 5.4 can hopefully be seen as a step towards a general methodology for situated reasoning.

Of course, there are many important issues which need to be investigated and clarified in fleshing out the methodology. In particular, the task of reasoning with background assumptions within the regularity itself has been replaced by the problem of choosing an appropriate channel to reason with. While this seems to be a more natural way to view the task, formulating this approach is not straightforward and it may turn out that the choice of channel is analogous to the problem of specifying a (globally) reliable regularity. However, I believe that the shift in approach will prove to be a useful one, although clearly much more work is required on this task.

Chapter 6

Summary and Further Work

This chapter contains a brief summary of the thesis and a discussion of possible future work. The future work involves both removing some of the technical shortcomings in the thesis (usually in the form of simplifications) as well as an outline of a plan for future research. The longer-term research aim is to use channel theory as a uniform formal framework for an integrated model of various aspects of "situated reasoning".

6.1 Summary of Thesis

This thesis has investigated the use of channel theory, Barwise and Seligman's recent mathematical model of information flow, as a formal semantic framework for several topics which I have collectively labelled *conditional reasoning.* I have shown how aspects of channel theory lead to analyses of conditional and generic sentences that address certain problems with traditional semantic approaches. The analysis of generics in particular benefits from the various different options offered by channel theory for representing exceptions to a general regularity.

My main concern throughout, however, has been to develop channel theoretic frameworks for reasoning with conditionals and generics, and for using generics to defeasibly reason about the properties of associated individuals. This led to a method for encoding background assumptions of regularities via a relation between channels and modifying the operations on channels to account for such assumptions. The resulting logics were compared to traditional logics of conditionals and default reasoning and shown to satisfy many important properties.

Channel Theory

Channel theory was developed by Barwise and Seligman as a formal theory of information flow in response to various problems with previous frameworks. In particular, channel theory addresses the problem as to how regularities between types can be both reliable and fallible. This problem has plagued philosophers for a long time and led to a view that the only "true" regularities must be tautological, since all background conditions of a regularity could never be fully specified. However, allowing a fallible relationship between types to still count as a *bona fide* regularity seems to avoid this problem.

The fallibility of regularities is modelled in channel theory via a two-level analysis, explicitly distinguishing connections at the token-level from links at the type-level. This allows for three different ways of modelling "error": by a token falling outside the scope of applicability of a regularity (i.e. not being in the token set of the associated classification); by a token not being involved in a token-level connection; by a token-level connection not being classified by a type-level constraint. I showed in Chapter 4 how these different sorts of "error" seem to account for different ways in which an individual can be an exception to a generic sentence (i.e. without invalidating the generic). A crucial property in the channel-theoretic account of error is that there is no need to explicitly represent the background assumptions of a regularity within the channel supporting that regularity—the fact that a token participates in a connection that is classified by the regularity reflects that "normal conditions" are in place.

Barwise and Seligman provide operations on channels that effectively lead to a "logic of channels", allowing a channel theoretic model of inference to be constructed. This logical framework has been my main focus throughout the thesis. I have taken initial observations by Barwise (1986) and ter Meulen (1986)—whereby they (respectively) claim that situation theoretic constraints could be used as a basis for a semantic analysis of conditional and generic sentences—and shown how logics of conditional and generic sentences (and for default reasoning) can be constructed using the channel operations. The view that fallible regularities—as they are modelled in channel theory—are the objects which such sentences describe fits well with the situation semantic view of natural language semantics.

Conditional Logics

The first major contribution of this thesis involved the development of a channel theoretic system of reasoning about conditionals. Barwise (1986) originally suggested that an account of the semantics of conditional statements could be given in terms of situation theoretic conditional constraints. Barwise and Seligman (1994) later suggested that a channel theoretic analysis would provide a conditional with demonstrative content, more in line with a situation semantic view of natural language semantics. In Chapter 3, I took the demonstrative content to be a channel and then focussed on the logic of conditionals that arose from the channel operations. The central issue in logics of conditionals is that Transitivity, Monotonicity and Contraposition must be invalidated. However, I showed that with the channel operations as they were defined by Seligman and Barwise (1993), this was not the case. The problem, I argued, was that the token-level of a channel *reflected* the underlying background conditions of a regularity but did not *capture* them.

To combat this problem I defined the notion of a channel hierarchy. This involves a *subchannel* relation between channels, allowing background assumptions of a regularity to be represented via the relationship between channels. The important aspect of this construct is that such assumptions are not represented explicitly within the channel $\mathcal{C}$ supporting the regularity, but only implicitly via $\mathcal{C}$'s relationship to other channels. This is critical since part of the motivation underlying channel theory is that such conditions are not explicitly represented and that any particular use of a channel may lead to an erroneous inference. Since reasoning with a channel does not involve the subchannel hierarchy but only the channel itself, this concern is respected.

To obtain a logic of conditionals that invalidated the above rules of inference, the standard channel operations were modified so as to take into account the background assumptions that were encoded via a channel hierarchy. The resulting logic did indeed invalidate Transitivity, Monotonicity and Contraposition, as was illustrated via various examples. An important issue that was then tackled in some detail involved determining which patterns of reasoning the system did validate. A number of axioms and rules of inference (taken from (Nute 1980)) were examined and, assuming certain reasonable conditions (some that simply imposed a certain "symmetry" on a hierarchy, others that any channel should be expected to satisfy), it was shown that many of these axioms and rules of inference were validated. In particular, the behaviour of the channel theoretic logic of conditionals had much in common with logics of conditionals based on relevant logic (which is also considered to be a "logic of information").

Logics of Generics

Recently, several researchers from the AI and philosophical communities have attempted to define logics of default reasoning starting with a conditional logic basis. Chapter 4 described a similar project within the channel theoretic framework. Several problematic issues in the semantics of generics were discussed and it was argued that they pointed to a need for a satisfactory analysis of generics to accommodate a notion of "context". It was proposed that the channel theoretic analysis of generics did just that, with a channel providing a context for the regularity which formed the descriptive content of a generic statement.

Various problematic examples for traditional analyses of generics were discussed and I proposed that the channel theoretic analysis, with its contextual aspect, gives a satisfactory account of these problems. The different ways by which channel theory accounts for error—via restriction to a particular set of tokens; via a token being a pseudo-signal; via a connection being an exception—seems to provide a comprehensive analysis of the various different ways in which an individual can be an exception to a generic without invalidating that generic. In particular, this seemed to give a more satisfactory analysis of certain examples that are particularly problematic for the "normative" approach to generics (such as the example involving peacocks that lay eggs and have bright feathers). Other features of the channel theoretic account, such as the non-reductionist account of regularities, meant that other problematic issues were also avoided.

Once again, my main concern lay in providing an adequate system of reasoning about generics, starting with the channel theoretic logic of conditionals as a basis. While the channel theoretic system for conditionals placed only minimal constraints on a channel hierarchy, I argued that a system for reasoning with generics required a tacit assumption that any background assumptions were implicit in the original collection of regularities. This was the *only* difference between the logic of conditionals and the logic of generics—i.e. there was no modification of the channel operations themselves, only in what was allowed as a channel hierarchy. This meant that the move from a conditional logic to a logic of generics did not involve what Morreau (1992a) called a "ghost in the machine". By defining a particular construction of a channel hierarchy (given an initial collection of channels), a specific system of reasoning with generics was obtained. It was shown that this system satisfies various important patterns of reasoning outlined in (Morreau 1992a), such as Specificity, Carlson's Graded Normality, and Addition of Generics. Various examples from the traditional AI literature on default reasoning were discussed and it was shown that the logic of generics satisfied the required behaviour for these examples.

Logics of Defeasible Reasoning

Morreau (1992a) argues that a critical issue in logics of generics is that they enable default reasoning about the properties of individuals. Under the channel theoretic account of generics of Chapter 4, generics are effectively "simply true", due to the non-reductionist account of regularities. Hence, a condition of "maximal normality" had to be imposed before the channel theoretic framework could be used to defeasibly reason about the individuals themselves. Such a condition was defined in Chapter 5. A default reasoning system, as defined in Chapter 5, consists of a hierarchy of information links, which is obtained directly from a hierarchy of channels comprising a system of generics, with the additional condition that it satisfies maximal normality.

As a way of evaluating the channel theoretic logic of default reasoning, Kraus et al.'s (1990) work on nonmonotonic consequence relations was discussed, and it was shown that under certain conditions, the channel theoretic default reasoning system defined a preferential consequence relation. Unfortunately, the question as to whether the channel theoretic default system defines a rational consequence relation was not answered, although it is conjectured that it does so.

The final section of Chapter 5 described how the maximal normality condition reflects a methodology towards reasoning with "approximate" regularities. Rather than reasoning with conditions that are assumed to hold by default, I argued that it was more natural to view such reasoning as being performed at a certain level in a hierarchy, with the appropriate level being determined using the available information. Such reasoning could, of course, lead to erroneous conclusions, which is inherent in the nature of default reasoning. However, this subtle difference seemed to encompass a more natural approach to "situated reasoning". The chapter concluded with an application of the methodology to AI's qualification problem, providing a much more natural solution than the usual one involving traditional default logics.

6.2 Remaining Issues and Future Research

This thesis is meant to be a step along the path of using a channel theoretic framework for various tasks in reasoning and semantic modelling. The specific tasks that have been addressed in this thesis, falling under the general topic of *conditional reasoning*, are logics of conditionals, generics and default reasoning. Several issues have been left unanswered in the treatments of these topics—in particular, each system can be generalised in certain ways. However, there are several related topics of interest that have been touched upon along the way and some of these promise to lead to interesting questions. In this section, I will summarise both the open questions of this thesis and the areas of potential interest that it has raised.

6.2.1 Tying Loose Ends

In various places in this thesis I have made certain simplifications, often in order to allow more direct comparison with standard "unsituated" logics of conditionals and default reasoning. The most general version of the channel theoretic framework does not involve such simplifications, of course, and work needs to be performed to investigate the properties of the channel theoretic logics in their more general settings. In particular, removing the simplifications of "decontextualisation" (e.g. ignoring the connection-level in the logic of generics; assuming all channels are reflexive; assuming there is only one classification involved in the logics of generics and default reasoning) will make the logics fit more comfortably with the methodological stance that reasoning is situated and dependent on context. My motivation for wanting to directly compare the channel theoretic systems to more standard logics was to demonstrate that the general framework (involving the encoding of background assumptions within a channel hierarchy) was able to model behaviour that was generally agreed to be desirable. However, having seen that the channel-theoretic framework can capture such behaviour in the restricted setting, exploration of the more general setting—and investigating the sorts of behaviour that the channel-theoretic framework offers and which *cannot* be adequately modelled with more traditional frameworks—becomes a higher priority.

Other ways in which the general channel theoretic framework could be generalised was discussed in Chapters 2, 3 and 4 and their respective appendices. One such task involved further investigating the way in which complex propositions (and propositions involving channels) can be represented within the general channel theoretic framework. Such objects should be expressible within the classification and channel theoretic frameworks, allowing a uniform representation of (among other things) nested conditionals and generics, and negated conditionals and generics (i.e. assertions that a channel does not support a conditional/generic). As discussed in the body of the thesis, there seems to be no inherent reason why this should not be possible. Another topic that requires further investigation involves incorporating more complex situation theoretic objects into the general framework. Again, there seems to be no reason why this should present any insurmountable problems, although the properties of such objects in the channel theoretic setting need to be explored. Incorporating situation theoretic objects would also allow more linguistically satisfying analyses of the semantics of conditional and generic sentences (including a more comprehensive modelling of the dyadic *GEN* operator).

Chapter 5 contains the most explicit "open problem" of the thesis—i.e. whether or not the channel theoretic model of default inference effectively defines a rational consequence relation. As I stated, I believe that the way in which a default channel hierarchy is constructed—whereby any background condition of a regularity must be present as a *defeater* in the initial collection of information—ensures this to be the case. However, further work on the representation of the negation of regularities (i.e. asserting that a regularity is *not* supported by a channel) needs to be performed before Rational Monotonicity can even be meaningfully expressed within the framework. The other slightly unsatisfactory issue of this particular chapter involved the definition of *external inappropriateness*. I am not convinced that this condition is actually necessary to prove the result of Theorem 5.3.2—it may be that the constrained nature of the construction of a default channel hierarchy negates the need for such a condition. This issue is not entirely clear at this stage.

6.2.2 Further Properties of the Framework

As well as addressing the issues described above, there are a number of areas of potential interest related to the channel theoretic framework in general, and the channel hierarchy framework in particular. I have already mentioned that an investigation of the properties of the situated or contextualised versions of the logics of conditionals and defaults would be an interesting topic, as would the general role of context in reasoning with conditionals and generics.

The introduction of a channel hierarchy will be seen as unfortunate by some—e.g. Restall (1994) and Seligman (personal communication) have each stated that they believe the use of channels as demonstrative content should invalidate the undesired rules of inference for conditionals without the need of any further mechanism. The discussion of Section 3.2.3 showed that this was not the case. However, the particular approach I have taken here is obviously not the only possible one, and the question remains as to whether a variant of this approach would lead to similar results via a more satisfying treatment.[1] It may be the case that the use of the channel hierarchy can be restricted to particular cases only (e.g. the motivating examples of Chapter 3 all involved a logical channel). The properties of the conditional operations themselves also bear further investigation. One of the particular questions raised in Chapter 3 regards whether it would not be more natural to restrict the Parallel Join operation rather than the Contraposition operation.[2]

A final topic of investigation involves the mathematical and logical properties of the channel theoretic model of information flow. I mentioned earlier that Girard's (1987) linear logic is considered to be a logic of information and similarities between channel theory and linear logic have been noted.[3] Another recent innovation that seems to have some similarities to channel theory is *arrow logic* (e.g. (van Benthem 1994)). Arrow logic is an abstract framework of dynamic logic involving the notion of an *arrow,* which is effectively a transition between states. Unlike in standard dynamic logic, an arrow is not in general identified with its source and target. Further, an arrow may support an implicational formula which basically describes how information at the source of the arrow allows inference regarding the target of that arrow. Under this simple description, there seems to be some (at least superficial) similarities between channel theory and arrow logic and it would be most interesting to investigate to what extent the former can be modelled by the latter,[4] especially since this may clarify the relationship between channel theory and other logics of information flow.

[1]Simple variants of the approach taken here include denying the existence of a composition channel or of the offending pair of the signalling relation—rather than the offending pair of the indicating relation—when conflict of background assumptions occurs.

[2]As discussed in that chapter, the most damning problem with Contraposition in a logic of conditionals can be alleviated by disallowing Weakening of the Consequent instead. However, Freund et al. (1991) argue that Contraposition is also an invalid pattern of inference for default reasoning.

[3]E.g. Jon Barwise and Vaughan Pratt have made such observations on the Linear Logic Mailing List. Barwise et al. (1994) have also defined a channel theoretic model of the Lambek calculus, which is closely related to the non-commutative implicational fragment of linear logic.

[4]It is unclear to me whether the concept of an exception could be modelled at all in arrow logic.

6.2.3 Multi-Agent Situated Reasoning

I talked a bit in Chapter 1 about the potential of channel theory as a framework for modelling multiple situated agents that interact with each other and their environment. Many of the important issues in this research area are closely related to the specific concerns addressed by situation theory, such as: the role of context and its effect on efficient reasoning; the effect of an agent's specific embodiment in the world on the content of its internal states; the importance of an agent's perspective or point of view on its environment; communication and collaboration between (possibly heterogeneous) agents. This section contains some speculative thoughts on how such issues may be addressed in a uniform channel theoretic framework.

The Role of Context in Reasoning

Chapter 1 contained a number of references to important pieces of work in which it was argued that accounting for the specific context an agent finds itself in is crucial to effective and efficient reasoning. Traditional AI approaches to the modelling of acting and reasoning do not model an agent reasoning "in a context" but rather ascribe a world-view to the agent. Recent work, however, has seen the development of a number of "logics of context" (e.g. (Buvač et al. 1995; Giunchiglia and Traverso 1993; Guha 1991; McCarthy 1993)). These frameworks add the notion of *context* to standard predicate calculus, where a context can be anything from situations of the situation calculus to separate logical theories (each context possibly involving a separate language). In these frameworks, context tends to be used to modularise reasoning systems. For example, Guha's (1991) *micro-theories* can be likened to modules of the full logical theory available to an agent. Specific reasoning tasks are performed within micro-theories designed for those tasks and information which is not available within the micro-theory for a task cannot be used on that task. The purpose of micro-theories is to focus and modularise reasoning and thereby make it more efficient.

An important facet of some of the logics of context is that information can be "lifted" and "lowered" between contexts (analogously to importing and exporting variables between modules in a programming language). Reasoning can be localised within a particular context with the resulting conclusions being lifted into another context which performs a separate (though related) reasoning task. This approach has been used in various models (particularly by McCarthy (1993) and Giunchiglia et al. (1993)) in models for reasoning about the beliefs of other agents.

The notion of "reasoning within a context" has clear similarities to the methodology for default reasoning sketched in Section 5.4, as demonstrated by Giunchiglia and Bouquet's (1994) approach to the qualification problem using a logic of context. One shortcoming of most of the logics of context of the AI literature is that most of the treatments are purely syntactic. Channel theory—which allows the contextualisation of information as well as the use of channels for exporting information between contexts—would seem to be a candidate for providing a corresponding semantics. It would also be worthwhile investigating the issues and approaches used by the proponents of the AI logics of context and seeing whether they cast any light on the channel theoretic methodology for situated default reasoning.

Perspective and Situated Content

A property that is closely related to context is the indexicality of reasoning, which involves accounting for the *perspective* or *viewpoint* of an agent on its environment. The indexicality of reasoning is pretty much ignored in the AI literature but—particularly in multi-agent settings—would seem to offer the possibility of more efficient systems. One would like to be able to specify or model the behaviour of each given agent in an environment from that particular agent's point of view, rather than needing to specify a *global* view of the environment which each agent then adopts. Agre and Chapman's (1987) *Pengi* system is the only AI system I know of that incorporates an explicit agent-centred view of the world. Lesperance and Levesque (1995) have recently proposed a modal logic of knowing and acting that incorporates an agent's particular point of view.

Modelling perspective and point of view within the channel theoretic framework should be a relatively straightforward process—in fact, classification theory seems naturally suited to this task. Let A and B be agents and for simplicity assume that they reason over the same sets of tokens and types.[5] We can associate classifications A' and B' with the respective agents, where A' and B' contain the same tokens and types. Now suppose A and B face each other across a table which contains a ball and a pen-

[5] This need not be the case, of course. In fact, Healey and Vogel (1994) take as a starting premise the assertion that *no* two agents can reason over the same tokens and types—i.e. types are somehow internalised to the agents and tokens are determined by the types they support.

cil. A' may support $(table : \ll leftOf, ball, pencil \gg)$ whereas B' may support $(table : \ll rightOf, ball, pencil \gg)$. Of course, these two propositions are related—since A and B face each other across the table then A''s supporting $(table : \ll leftOf, ball, pencil \gg)$ carries the information that B' supports $(table : \ll rightOf, ball, pencil \gg)$ and vice versa. This information is carried in a channel $\mathcal{C} : A' \Rightarrow B'$. If A is attuned to this channel, then she can infer something about B's mental state (although this inference may be incorrect and may therefore correspond to an exception in $\mathcal{C}$).

Related (in fact, in some way generalising) indexicality is the situatedness of the content of internal states, an issue which has received much attention in the philosophical literature. Dretske (1988) shows how environmental constraints—particularly those linking an agent's internal states to its embedding environment—can be exploited in the design of simple systems that behave in seemingly complex ways in a rich environment. To take the *frog* example from Chapter 2, the frog seems to be feeding on flies because it is hungry, even though it cannot help itself since darting its tongue at "black flying things" in its immediate vicinity is a reflex action.

Rosenschein and Kaelbling's work on *situated automata* (Rosenschein 1985; Rosenschein and Kaelbling 1986) is perhaps the most notable AI attempt to explicitly design machines based on this principle and provide a logical framework for explaining their behaviour. A situated automaton is only a very simple finite state machine, but each of its states is related to a particular condition in the world. The workings of the state transition function can then be explained using intentional language—i.e. by describing the behaviour in terms of the outside environment and the machines "beliefs" about this environment.

The regularity between an agent's internal state and its external environment is just one particular sort of regularity[6] and hence can be modelled within the channel theoretic framework. Channels supporting such regularities can be used to infer the content of a simple agent's internal states (with respect to some model of the world) and can thereby be used to ascribe complex intentional behaviour to the agent (within the theorist's model) without requiring the agent itself to possess the correspondingly rich discriminatory and complex processing powers—i.e. the complex intentional view is *only* available to the theorist because of the perceived regularities between the (simple)

[6]Phil Agre (personal communication) considers it the most important sort of regularity when modelling situated agents.

internal states of the agent and the complex (i.e. rich in types) environment. The resulting model involves a network of channels—some representing the link between the agent and the world, the rest representing other regularities in the world—with the simple agent and its associated deficient classification sitting in the middle. This would allow the logical design and specification of simple agents for performing (seemingly) complex reasoning tasks.

Interacting Heterogeneous Agents

I have outlined above how the notion of "perspective" and "content of mental state" can be modelled in a channel theoretic setting. When multiple agents are involved, not only does each of these characteristics have to be modelled for each agent, but the interaction between the agents must also be modelled. This interaction is usually with respect to some "external" medium. If we assume agents to be heterogeneous in their outlook on the world, then this external medium cannot be some "objective" world unless it is the theorist's own view of the world.[7] One way in which an objective (or at least *perceived* to be objective) medium is provided is via communication, particularly linguistic communication.

Healey and Vogel's (1994) channel theoretic model of dialogue addresses the problems of agents that interact linguistically on some task for which the different agents may disagree on certain ontological issues. In particular, this may lead to communication breakdown even though each dialogue participant considers each step of the dialogue to involve successful negotiation. Healey and Vogel show how the channel theoretic concepts of pseudo-signal and exception can be used to explain various sorts of communication breakdown that occur in dialogues involving negotiation. Their model involves each agent being represented by a distinct collection of classifications (with no tokens or types being shared across agents), with the only interaction between the agents being provided by a single *utterance situation* which is shared between the agents.[8] An analogous model could be used to model multiple interacting agents, with each agent

[7] This is, of course, a useful stance to take—i.e. that the theorist discriminates some objective world in which the agents operate—particularly when using the framework for the specification of a multi-agent approach to solving some task.

[8] Personally, I would model the task slightly differently to the way Healey and Vogel do. In particular, I would prefer an explicit *communication channel* linking private utterance sites. The fact that the participants classify the utterance site in the same way can only be established by an external observer and it seems more natural (to me) to instead introduce a single "external" channel.

being provided with its own frame of reference and reasoning. The interaction between agents would be confined to some interface, which is modelled by channels specifying the information-flow between the classifications representing the different agents. Such a model would lead to a methodology for the specification of multi-agent systems—each agent could be specified independently of the others, with the total system being specified on definition of the various agents' interaction.

The dialogue model of Healey and Vogel also seems to offer a framework for a potentially interesting extension of some of the work described in this thesis. In the channel theoretic semantic model of generics, the demonstrative content of a generic is a channel. Exactly *which* channel is the intended one is not always clear from the context and is certainly not specified by the theory itself. It may well be that a speaker of an utterance (involving a generic sentence) intends a particular channel as the demonstrative content while the hearer understands the generic as being about a different channel. Since channels form the basis of inference about individual tokens (in the model of default reasoning of Chapter 5), this may lead to erroneous conclusions about individuals. For example, suppose A utters

"Dutch make good sailors."

meaning that the channel $\mathcal{C}$ which contains only *sailor* tokens supports the regularity $dutch \rightarrow good.sailor$—i.e. what A means to convey is that Dutch sailors make good sailors. The hearer B may take this utterance to be about a channel $\mathcal{C}'$ containing all Dutch people—i.e. B takes the utterance to mean that *all* Dutch people make good sailors. If the above assertion is made in a context where A is indicating to B as to who he should recruit for his navy, then B may well fill his boats with Dutch farmers, leading to some problematic sea voyages! The Healey/Vogel model of dialogue seems to be a natural framework in which to model dialogues involving generics and the sort of erroneous "default information" that may be inferred when miscommunication of the above sort occurs.[9]

[9]Sheila Glasbey and I have commenced investigation into modelling such phenomena.

Appendix A

Channel Theory: Further Formal Definitions

In this appendix, I elaborate on the formal concepts defined in Chapter 2.

A.1 Simple Situation Theoretic Universes

This section contains a more precise and complete presentation of Barwise and Cooper's (1991) formal definition of the concepts of situation theory. As they acknowledge, their definition makes use of a great deal of work of many others. The definition below is borrowed from (Barwise and Cooper 1991) (with a few minor cosmetic changes).

Definition

A *simple situation theoretic universe* consists of various sorts of objects and operations on them, satisfying the conditions listed below.

Objects

1. Infons;

2. Propositions, some of which are true, the rest false;

3. Situations;

4. Abstracts, including:
   - infon abstracts, which are called *relations;*
   - proposition abstracts, which are called *types,* and which include the following:
     - types *Infon, Proposition, Abstract, Relation, Type, Situation* and *Set,* corresponding to the sorts of objects;
     - a type $\models$;
     - a type *IsOf;*
     - the collection of infons;
5. Parametric objects corresponding to each of the above;
6. Parameters;

   Note: an *indexed family F of parameters* is a one-one function whose range is a set of parameters. The domain of F is called the set of *indices* of F. A *finite sequence* $[X_1, ..., X_n]$ *of parameters* is an indexed family of parameters whose index set is the set $\{1, ..., n\}$. An *anchor* is a function h whose domain is a set of parameters.
7. Sets.

   Note: the inclusion of set theory allows the representation of functions. Functions will often be referred to as *assignments,* since that is how they will be used.

Operations

We define operations on the sorts of objects in a universe, as follows.

Infon operations:

1. Two partial binary operations $\ll . \gg^+$ and $\ll . \gg^-$, whose first argument is a relation r and whose second argument is an assignment f, with the results (if defined) being infons. We write $\ll r, f; 1 \gg$ and $\ll r, f; 0 \gg$ for $\ll r, f \gg^+$ and $\ll r, f \gg^-$ respectively.
2. Two total binary operations $\wedge$ and $\vee$ from infons to infons.

Proposition operations:

3. A partial binary operation $(. : .)$ whose first argument is an assignment f and second argument is a type T, with the result (if defined) being a proposition $(f : T)$.

4. Two total binary operations $\wedge$ and $\vee$ from propositions to propositions.

5. One total unary operation $\neg$ from propositions to propositions.

Operations on parametric objects:

6. An operation *par* which assigns to each object o a set $par(o)$ of parameters, called the set of parameters of o.

7. A partial binary operation $.[.]$ of substitution whose first argument is an object o and second argument an anchor f.

8. A binary operation λ whose first argument is an indexed family F of parameters and whose second argument is a parametric object o. The operation is defined for all such arguments and the result $\lambda F o$ is an abstract. If o is an infon, then $\lambda F o$ is called an *infon abstract,* and similarly for the other sorts of objects.

Restricting objects:

9. A partial binary operation $. \downarrow .$ whose first argument o is an arbitrary object and whose second argument is a (possibly parametric) proposition.

Operations on abstracts:

10. An operation *Ext* which assigns to each non-parametric type T its *extension* $Ext(T)$, a collection of (possibly parametric) assignments.

11. An operation *Approp* which assigns to each abstract b a type of assignment $Approp(b)$ (i.e. $Ext(Approp(b))$ is a collection of assignments). $Approp(b)$ is called the *appropriateness conditions* for b. An assignment f is *appropriate for* a non-parametric abstract b if $f \in Ext(Approp(b))$.

12. A binary operation *Apply* whose first argument is an abstract b and whose second argument is an assignment f.

Axioms

To qualify as a simple situation theoretic universe, the objects and operations must satisfy the following axioms. In what follows, = is Kleene equality: "b = c" means that "b" and "c" either both denote or neither does, and if they denote, then they denote the same object.

Parameters:

1. $par(o)$ satisfies the following conditions:
   - $par(o) = \emptyset$ iff o is non-parametric;
   - $par(X) = \{X\}$ if X is a parameter;
   - $par(\ll r, f; i \gg) = par(r) \cup par(f)$;
   - $par(\sigma \wedge \tau) = par(\sigma) \cup par(\tau)$, and similarly for $\vee$;
   - $par((f : T)) = par(f) \cup par(T)$;
   - $par(p \wedge q) = par(p) \cup par(q)$, and similarly for $\vee$ and $\neg$;
   - $par(o \downarrow p) = par(o) \cup par(p)$;
   - $par(Approp(b)) \subseteq par(b)$;
   - $par(\lambda F o) = par(o) - rng(F)$, for F an indexed family of parameters;
   - $par(o[f])$ is the set of parameters in o, less those in $dom(f)$, together with all the parameters in any $f(X)$, for $X \in par(o)$.[1]

Appropriateness conditions and extensions:

2. There is a type T_ϕ for every formula $\phi(x_1, ..., x_n)$ of set theory with atoms, for which every assignment $[a_1, ..., a_n]$ is appropriate. The extension of T_ϕ consists of those assignments $[a_1, ..., a_n]$ that satisfy ϕ when x_k is interpreted by a_k for each k.

3. The appropriateness conditions of an abstract b determine when b combines with other objects, either by substitution or (in the case of relations and types) by predication. More precisely:

   Given an indexed family F and an assignment f to the domain of F, $\hat{f}$ is the anchor for the parameters in $rng(F)$ with $\hat{f}(X) = f(F^{-1}(X))$—i.e. $\hat{f}$

[1] I.e. $par(o[f]) = (par(o) - dom(f)) \cup \{par(f(X)) \mid X \in par(o)\}$.

anchors X with the object assigned to the role index corresponding to X in F.

(a) If $c = \lambda F o$ then an assignment f is appropriate for c iff $o[\hat{f}]$ is defined.

(b) If $T = \lambda F p$ is a non-parametric type, then for any assignment f appropriate to T, $f \in Ext(T)$ iff $p[\hat{f}]$ is true.

(c) For any non-parametric relation r, an assignment f is appropriate for r iff $\ll r, f; 1 \gg$ and $\ll r, f; 0 \gg$ are defined. Similarly, for any non-parametric type T, f is appropriate for T iff $(f : T)$ is defined.

(d) An assignment f is appropriate for the basic types *Infon, Proposition, Abstract, Relation, Type, Situation* and *Set* just in case f is unary, say $f = [b]$. In this case, f is of one of these types just in case b is a non-parametric infon, proposition, abstract, relation, type, situation or set, respectively.

(e) The appropriate assignments for $\models$ are the assignments $[s, \sigma]$, where s is a (possibly parametric) situation and σ is a (possibly parametric) infon. (We write $s \models \sigma$ rather than $[s, \sigma] \in Ext(\models)$ for all appropriate assignments to $\models$.)[2]

   i. $s \models \ll \lambda F \tau, f; 1 \gg$ iff $s \models \tau[\hat{f}]$.
   ii. If $s \models \ll \lambda F \tau, f; 0 \gg$ and s is coherent then $s \not\models \tau[\hat{f}]$.

(f) The appropriate assignments for $IsOf$ are the assignments $[a, T]$ such that T is a type and a is appropriate for T.

(g) If $[a, T]$ is non-parametric and appropriate for $IsOf$, then $[a, T] \in Ext(IsOf)$ iff $a \in Ext(T)$. (We write $a : T$ rather than $[a, T] \in Ext(IsOf)$. Note that if a is appropriate for T then $(a : T)$ is a proposition such that $a : T$ iff $(a : T)$ is true (see below)).

   The appropriate assignments for a non-parametric infon σ are just those unary assignments s such that s is a non-parametric situation, and for such an assignment, $s \in Ext(\sigma)$ iff $s \models \sigma$.

(h) If o is an object and p a proposition, then

   i. $o \downarrow p$ is defined iff p is not a false proposition.
   ii. If p is true (and hence non-parametric), then $o \downarrow p = o$.

[2]I do not include the first two conditions of Barwise and Cooper's definition, involving the conjunction and disjunction operations, for reasons discussed below.

iii. $. \downarrow .$ distributes over the various closure operations that do not abstract over or substitute for parameters; e.g. $(\sigma \wedge \tau) \downarrow p = \sigma \downarrow p \wedge \tau \downarrow p$.

iv. $\lambda F o \downarrow p = \lambda F(o \downarrow p)$ if no parameters in p are in the range of F.

Truth:

Note: The following theory of truth applies to non-parametric propositions only.

4. $(f : T)$ is true iff $f \in Ext(T)$.

5. $p \wedge q$ is true iff p and q are both true. Similarly, $p \vee q$ is true iff either p or q is true.

6. $\neg p$ is true iff p is false.

Abstraction, application and substitution:

7. $X[f] = f(X)$ if $X \in dom(f)$ and $X[f] = f(X)$ if $X \notin dom(f)$, for any parameter X.

8. $o[f] = o$ if f is the empty function.

9. $o[f] = o[f_0]$, where f_0 is the restriction of f to $par(o)$.

10. .[.] distributes over the various infon and propositional operations; e.g.

$$\ll r, a; i \gg[f] = \ll r[f], a[f]; i \gg$$
$$p \wedge q[f] = p[f] \wedge q[f]$$
$$o \downarrow p[f] = o[f] \downarrow p[f]$$
$$\lambda F o[f] = \lambda F(o[f]) \text{ provided } dom(f) \text{ is disjoint from } rng(F).$$

11. If $dom(F)$ is disjoint from $par(\lambda G o')$ then $\lambda F o = \lambda G o'$ iff $dom(F) = dom(G)$ and o is the result of simultaneously substituting $F(G^{-1}(X))$ for X in o', for all parameters $X \in rng(G)$.[3]

12. $Apply(\lambda F o, f)$ is the result of substituting $f(i)$ for the parameter X_i in o, provided f is an assignment that is appropriate for b. More precisely:

    - $Apply(\lambda F o, f)$ is defined iff $f \in Ext(Approp(\lambda F o))$.
    - For any such f, $Apply(\lambda F o, f) = o[\hat{f}]$.

[3]This axiom effectively allows replacement of parameters.

Structure axioms:

The following axioms are taken to be optional by Barwise and Cooper. However, I include them so as to ensure the fine-grained intensionality I desire.

13. Basic infons and propositions are uniquely determined by the constituents out of which they are formed. More precisely:

    - If $(a : T) = (a' : T')$ then $a = a'$ and $T = T'$.
    - If $\ll r, a; i \gg = \ll r', a'; j \gg$ then $r = r'$, $a = a'$ and $i = j$.

Barwise and Cooper propose further possible axioms that could be added to the ones above. I will not present such axioms here (although some of them deal with important concepts of situation theory, such as the *part-of* relation) since they play no part in the results of this thesis. Note that, when I define a situation theoretic classification (i.e. tokens are situations and types are situation-types), then the objects of this classification are assumed to satisfy all the general conditions of classifications, as set out in Section 2.2. At the time of writing, no model of the Situation Theoretic Universe axioms has been defined, so there is no guarantee that they are consistent; however, Barwise and Cooper believe that they are so.

The only difference of any significance I have made to Barwise and Cooper's definition is in the axiom involving the extension of the type $\models$. Barwise and Cooper include the following two conditions, which I omit:

i. $s \models \sigma \wedge \tau$ iff $s \models \sigma$ and $s \models \tau$;

ii. $s \models \sigma \vee \tau$ iff $s \models \sigma$ or $s \models \tau$.

On the one hand, I have no need for these conditions, since the interaction of the $\models$ type with the infon-conjunction and -disjunction operations is handled by the conditions I impose on classifications with type-structure, as presented in Section 2.2.2. However, even there I do not impose a condition that captures the second of those above, i.e. the condition involving the disjunction operation. In traditional situation semantics (e.g. (Barwise and Perry 1983)), a situation that supports a disjunctive item of information must support one of the disjuncts. However, in traditional logic, a disjunctive sentence can be valid without either of its disjuncts being valid, and since this thesis involves modelling logical frameworks with this behaviour, I wish to leave this option open, in general.

A.2 Operations on Classifications

In Section 2.2.3, I presented the definition of the sequential conjunction operation on classifications, one of several operations defined in (Seligman and Barwise 1993). Although I have not made use of such operations in this thesis (the assumption that all "logical" structure is represented at the type-level means that they are not needed), the framework could be significantly generalised by doing so. In particular, a more "contextualised" version of the framework could be defined, whereby, for instance, negations and conjunctions of Austinian propositions involving different tokens could be represented.

Some of the concepts defined earlier can be extended to take advantage of a framework generalised in this way. For example, the *subclassification* relation could be extended so that, for classifications A and B, A is a subclassification of $A \otimes B$ and $B \otimes A$,[4] leading to an algebra of classifications.

Such operations could also be extended to channels leading to the notion of "merging" information flow. In fact, one may be able to replace some of the current channel operations defined in Section 2.3.7 by classification operations. Alternatively, the channel operations themselves could be generalised by using the operations on classifications. For example, the current parallel meet (and parallel join) composition operation requires that the two channels to be composed link the same classifications. However, this restriction could be removed by defining parallel meet as follows.

Definition *Let $\mathcal{C}_1 : A_1 \Rightarrow B_1$ and $\mathcal{C}_2 : A_1 \Rightarrow B_1$ be channels.*

The definition of a connection graph is modified so that $\langle \wedge, a_1, a_2 \rangle \overset{(c_1 \| c_2)}{\mapsto} \langle \wedge, b_1, b_2 \rangle$ iff $a_1 \overset{c_1}{\mapsto} b_1$ and $a_2 \overset{c_2}{\mapsto} b_2$.

The parallel meet *of $\mathcal{C}_1$ and $\mathcal{C}_2$ is a channel $(\mathcal{C}_1 \triangle \mathcal{C}_2) : (A_1 \otimes A_1) \Rightarrow (B_1 \otimes B_2)$ defined as follows:*

1. $tok(\mathcal{C}_1 \triangle \mathcal{C}_2) = \{(c_1 \| c_2) \mid c_1 \in tok(\mathcal{C}_1), c_2 \in tok(\mathcal{C}_2)\}$;
2. $typ(\mathcal{C}_1 \triangle \mathcal{C}_2) = \{\langle \otimes, \gamma_1, \gamma_2 \rangle \mid \gamma_1 \in typ(\mathcal{C}_1)$ *and* $\gamma_2 \in typ(\mathcal{C}_2)\}$;
3. $(c_1 \| c_2) :^+ \langle \otimes, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \triangle \mathcal{C}_2)$ *iff* $c_1 :^+ \gamma_1$ *in* $\mathcal{C}_1$ *and* $c_2 :^+ \gamma_2$ *in* $\mathcal{C}_2$;

   $(c_1 \| c_2) :^- \langle \otimes, \gamma_1, \gamma_2 \rangle$ *in* $(\mathcal{C}_1 \triangle \mathcal{C}_2)$ *iff* $c_1 :^- \gamma_1$ *in* $\mathcal{C}_1$ *or* $c_2 :^- \gamma_2$ *in* $\mathcal{C}_2$;

[4] Similarly, $A \oplus B$ and $B \oplus A$ would be subclassifications of A, where $\oplus$ is a suitably defined *sequential disjunction* operation on classifications.

4. $source_{(\mathcal{C}_1 \,\triangle\, \mathcal{C}_2)}(c_1 \parallel c_2) = \langle \wedge\, , source_{\mathcal{C}_1}(c_1), source_{\mathcal{C}_2}(c_2)\rangle$;

   $target_{(\mathcal{C}_1 \,\triangle\, \mathcal{C}_2)}(c_1 \parallel c_2) = \langle \wedge\, , target_{\mathcal{C}_1}(c_1), target_{\mathcal{C}_2}(c_2)\rangle$;

5. $ante_{(\mathcal{C}_1 \,\triangle\, \mathcal{C}_2)}(\langle \otimes, \gamma_1, \gamma_2\rangle) = \langle \wedge\, , ante_{\mathcal{C}_1}(\gamma_1), ante_{\mathcal{C}_2}(\gamma_2)\rangle$;

   $succ_{(\mathcal{C}_1 \,\triangle\, \mathcal{C}_2)}(\langle \otimes, \gamma_1, \gamma_2\rangle) = \langle \wedge\, , succ_{\mathcal{C}_1}(\gamma_1), succ_{\mathcal{C}_2}(\gamma_2)\rangle$.

This generalised version of parallel meet composition provides a better model of the merging of the audio and visual channels of the television example of Section 2.3.7—the audio channel $\mathcal{A}$ and visual channel $\mathcal{V}$ involve different classifications but can still be merged to provide an all-round sensory information channel.

A.3 Parametric Types in Channel Theory

In this section, I describe a proposal for a possible treatment of parameters within the channel theoretic. The motivation behind the proposed approach is that (i) it should require minimal modification to the channel theoretic framework, and (ii) the properties of the resulting framework should be analogous to those of other semantic fragments—in particular, the parameters should behave similarly to discourse referents in DRT (Kamp 1981; Kamp and Reyle 1993) and variables in the generic operator *GEN* in that variables in the antecedent are effectively universally quantified (or generically quantified) while variables in the succedent (and not appearing in the antecedent) are existentially quantified.

A.3.1 Classifications with Parametric Types

Any classification A' containing parametric situation-types is based on a classification A involving only non-parametric situation-types, in the following way.[5] In the following, I tend to take an anchor as assigning objects directly to parameters rather than to roles. This is simply a notational convenience that simplifies the presentation.

Definition *Let A be a classification containing only non-parametric situation-types. The* classification $A[X]$ based on A and X, *where X is a set of parameters, is such that*

[5]The situation-types of A could contain parameters but any such parameters will be bound by the abstraction operation. For example, the situation-type $\lambda X[X \models \ll run, john; 1 \gg]$ contains the parameter X but X is bound by the abstraction operation. What counts, of course, is that $par(\phi) = \emptyset$ for any such type ϕ.

$tok(A[X]) = tok(A)$ and $typ(A[X]) = \{\phi' \mid \phi \in typ(A)$ and there is an anchor f such that $dom(f) \subseteq X$ and $\phi'[f] = \phi\}$. I require $(s : \phi)$ to hold in $A[X]$ iff there exists an anchor f such that $(s : \phi[f])$ holds in in A. As an abbreviation, I will sometimes write $s :^f \phi$ in $A[X]$ when $a : \phi[f]$ in A.

The notion of a "classification $A[X]$ being based on classification A" could be made more rigorous by employing Aczel and Lunnon's (1991) notion of a *universe with parameters.* Aczel and Lunnon define a universe with parameters $\mathcal{A}$ (which is basically a universe of parametric objects) as being *generated* from a base universe $\mathcal{A}_{np}$, where $\mathcal{A}_{np}$ contains only non-parametric objects. They show that, given $\mathcal{A}_{np}$ and a class X of parameters, the universe with parameters $\mathcal{A}$ generated is unique up to isomorphism. By taking $typ(A)$ in the definition above to be a universe[6] in the Aczel-Lunnon sense, we can simply take $typ(A[X])$ to be the universe with parameters generated from $typ(A)$ and X.

Of course, even having more rigorously defined the notion of a classification with parametric types, this still leaves freedom of choice as to how the support of parametric situation-types relates to non-parametric ones. Many authors are agnostic about such matters—e.g. Barwise and Cooper (1991) avoid the issue of truth for propositions involving parametric types altogether. The above definition effectively existentially quantifies over parameters, which I take to be a reasonable choice—a parametric situation-type ϕ is supported by a situation s exactly when the parameters in ϕ can be grounded in such a way that the corresponding instantiated situation-type is supported by s. For example, s supports the parametric type $\ll runs, X; 1 \gg$ exactly if there is some object o such that s supports $\ll runs, o; 1 \gg$—i.e. if there is *some* object that is running in s. The more interesting questions arise when such classifications are involved in a channel.

A.3.2 Channels and Parametric Types

I now turn to the properties of channels linking classifications whose types may be parametric. The model that I present below will be seen to satisfy the following criteria:

1. A channel is (basically) a classification and should be treated exactly as any other classification. The model below does just this, with the slight addition of an axiom

[6]Actually, $typ(A)$ will be the basis of a universe—there is a bit more to being a universe than being a set of structured objects.

that channels are assumed to satisfy. Also, the definition of a channel should not be modified—i.e. a channel involving parametric types should be a special case both of standard channels and of a classification involving parametric types.

2. The basic behaviour I require from constraints containing parameters is that any parameter in the antecedent is effectively universally quantified, while any remaining parameter in the succedent is effectively existentially quantified. For example, a constraint of the form $p(X) \rightarrow q(X, Y)$ behaves analogously to the predicate logic formula $\forall X(p(X) \supset \exists Y q(X, Y))$. This treatment is in line with the treatment of discourse referents in DRT (Kamp 1981; Kamp and Reyle 1993) and in the generic operator *GEN*.[7]

3. Since parameters are first-class objects, they should be effectively "propagated upwards" when we have nested constraints. For example, given a channel $\mathcal{C} : A \Rightarrow \mathcal{C}'$, where $\mathcal{C}'$ is itself a channel, then for any constraint γ in $typ(\mathcal{C})$ containing parameter X in both its antecedent and succedent, anchoring X in $ante(\gamma)$ should effectively simultaneously anchor X in $ante(succ(\gamma))$ and $succ(succ(\gamma))$.

For the first part, I assume that channels are not nested—i.e. any classification linked by a channel $\mathcal{C}$ is not itself a channel. This restriction is made for the sake of simplicity and is relaxed below. Let $A[X]$ and $B[X]$ be classifications based on A and B in the way described in Section A.3.1. A channel $\mathcal{C} : A[X] \Rightarrow B[X]$ is a triple $\langle f_L, C, f_R \rangle$, defined in the usual way. The first criterion above is achieved by ensuring that types in channels are never parametric. This is not to say that *constraints* do not contain parameters—this is the very thing I want to model so as to capture regularities on a more abstract level of generalisation. Instead, there is a distinction between the types of a channel and the antecedent and consequent types of the associated constraint—the latter are determined through use of the homomorphisms attached to the channel and may well contain parameters. For example, suppose the channel $\mathcal{C} : A[X] \Rightarrow B[X]$ contains the (non-parametric) type γ; this type can still capture the more general regularity associated with the use of parameters if f_L and/or f_R map γ to types in $A[X]$ and $B[X]$ containing (free) parameters.

Since channel-types are parameter-free, the concept of anchoring them is rather pointless (although still well-defined, of course). However, I introduce the following

[7] This is also the condition imposed on parametric constraints by Nivre's Axiom 7.2 (Nivre 1992, p. 136).

notion of the "anchored variant" of a channel-type γ, which is basically obtained by anchoring the parameters in $ante(\gamma)$ and $succ(\gamma)$.

Definition *Let $\mathcal{C} : A[X] \Rightarrow B[X]$ be a channel with $\gamma \in typ(\mathcal{C})$. The constraint σ is an* anchored variant *of γ* (wrt f) *in $\mathcal{C}$, written $\sigma = (\gamma \downarrow f)$, if $\sigma \in typ(\mathcal{C})$, $ante(\sigma) = (ante(\gamma))[f]$ and $succ(\sigma) = (succ(\gamma))[f]$.*

While the definition of a channel is exactly the usual one (in line with the first criterion above), I impose the following axiom on channels involving parametric types. (A *ground anchor* is one whose range contains only non-parametric objects.)

> **Parametric Constraint Axiom:** For any channel $\mathcal{C} : A[X] \Rightarrow B[X]$, type $\gamma \in typ(\mathcal{C})$ and appropriate ground anchor f such that $dom(f) = par(ante(\gamma))$ and $ante(\gamma)[f] \in typ(A)$, we require that
>
> *1.* $(\gamma \downarrow f) \in typ(\mathcal{C})$, and
>
> *2.* if $(c : \gamma)$ holds in $\mathcal{C}$ then $(c : (\gamma \downarrow f))$ holds in $\mathcal{C}$.

The first part of this axiom requires that, for a given constraint $\phi \rightarrow \psi$ in $\mathcal{C}$, all instantiations of the constraint (obtained by appropriately grounding the parameters in ϕ) are themselves constraints in $\mathcal{C}$. The second part ensures that if a constraint holds of a connection c then every instantiated instance of it also holds of c. This is closely related to the condition imposed by Nivre's Axiom 7.1 (Nivre 1992, p. 136), although Nivre uses the notion of constraints being factual (i.e. supported by some situation) rather than classifying connections.

I now turn to the question of what constitutes an exception to a parametric constraint. For a constraint linking parametric types, it is not so clear as to what should correspond to an exception. A connection c may be an exception to some instances of the constraint while satisfying other instances. For example, given a constraint such as $bird(x) \rightarrow flies(x)$, a connection $s \mapsto s'$ may be an exception to *opus* being a bird and flying while supporting the constraint for *tweety*. The approach I take here is that a connection c is an exception to a constraint if c is an exception to *any* instance of that constraint. In fact, this is captured by the most obvious definition of exception.

Definition *Given a channel $\mathcal{C} : A[X] \Rightarrow B[X]$, a connection c is an* exception *to the constraint γ iff $(source(c) : ante(\gamma))$ holds in $A[X]$ but it is not the case that $(c : \gamma)$ holds in $\mathcal{C}$.*

This, of course, is simply the standard definition for exceptions, as it is defined for the parameter-free case. The following result shows that it results in the behaviour that I have claimed for it.

Proposition A.3.1 *Let $\gamma \in typ(\mathcal{C})$ and $c \in tok(\mathcal{C})$ be such that c is an exception to $(\gamma \downarrow f)$ for some ground anchor f (where $dom(f) = ante(\gamma)$). Then c is an exception to γ.*

Proof Since c is an exception to $(\gamma \downarrow f)$, then $source(c) :^{f} ante(\gamma)$; hence, $source(c) : ante(\gamma)$ by definition. Further, it is not the case that $(c : (\gamma \downarrow f))$ holds, which means, by the parametric constraint axiom, that it is also not the case that $(c : \gamma)$ holds. □

Note that the converse does not necessarily hold—c can be an exception to some constraint without necessarily being an exception to any instance of it. Again, I would argue that this is the way it should be—just because each known instance of the constraint classifies a connection, it is no guarantee that the parametrised constraint also classifies it.

It is worth exploring this behaviour in a little more detail, especially with respect to how this relates to generic relationships. Consider a generic such as "Birds fly". In the model described in Chapter 4, this generic describes a channel $\mathcal{C}$ containing a constraint $bird \rightarrow flies$, with connections between bird-tokens. However, if we instead model the generic by a channel $\mathcal{C}'$ containing the parametric constraint $bird(X) \rightarrow flies(X)$, with connections between situation-tokens, then a connection $s \mapsto s$ is classified by the constraint only if *every* way of anchoring X (in such a way that $bird(X)[f]$ is a type in the antecedent classification) results in an instantiated constraint that also classifies $s \mapsto s$—i.e. if *every* (appropriate) bird flies in s. While this may at first seem anomalous, it simply means that $\mathcal{C}$ and $\mathcal{C}'$ correspond to different generics—$\mathcal{C}$ models the generic "Birds fly", while $\mathcal{C}'$ models a generic of the form "In (normal) situations, all birds fly".

I now consider the other criteria spelt out above, namely, that involving the implicit quantification of parameters and that involving nested channels. The satisfaction of the first of these is demonstrated by the following result.

Proposition A.3.2 *Let $\mathcal{C} : A[X] \Rightarrow B[X]$ be a channel with $c \in tok(\mathcal{C})$ and $\gamma \in typ(\mathcal{C})$, where $source(c) = a$, $target(c) = b$, $ante(\gamma) = \phi$ and $succ(\gamma) = \psi$.*

1. *If* $(c : \gamma)$ *holds then* $a : \phi$ *and* $b : \psi$*;*
2. *For any appropriate anchor* f*, if* $(a : \phi[f])$ *holds in* A *and* $(c : (\gamma \downarrow f))$ *holds in* $\mathcal{C}$ *then there exists an anchor* f' *such that* $(b : \psi[f][f'])$ *holds in* B*.*

Proof The first result is easily shown. To show the second, suppose the antecedent condition holds. Then $(b : \psi[f])$ holds in $B[X]$, which, by definition, implies the result. $\square$

Finally, I want to allow the possibility of a channel being nested. The desired behaviour is achieved by extending the notion of anchored variant as follows, ensuring that instantiations are propagated downwards into the nested constraints.

Definition *Let* $\mathcal{C}$ *be a classification (possibly a channel), with* $\gamma \in typ(\mathcal{C})$*. The* anchored variant of γ wrt an anchor f in $\mathcal{C}$*, written* $(\gamma \downarrow_{\mathcal{C}} f)$*, is defined as follows:*

- *If* $\mathcal{C}$ *is not a channel, then* $(\gamma \downarrow_{\mathcal{C}} f) = \gamma[f]$*;*
- *If* $\mathcal{C} : A \Rightarrow B$ *is a channel (*A *and* B *are possibly channels), then* $(\gamma \downarrow_{\mathcal{C}} f)$ *is a type* $\sigma \in typ(\mathcal{C})$*, if such a type exists, such that* $ante(\sigma) = (ante(\gamma) \downarrow_A f)$ *and* $succ(\sigma) = (succ(\gamma) \downarrow_B f)$*;*
  if no such type σ *fitting this description exists, then* $(\gamma \downarrow_{\mathcal{C}} f)$ *is undefined.*

By modifying the Parametric Constraint Axiom so as to make use of this revised notion of anchored variant, the desired behaviour for nested parametric constraints is obtained.

A.4 Further Formal Properties of Classifications and Channels

This section contains some further results regarding the formal properties of channels.

A.4.1 Further Properties of the Composition Operations

The following results prove the various properties claimed of the channel operations in Section 2.3.7.

The first property demonstrates that the objects defined by the parallel composition operations are themselves channels.[8]

Proposition A.4.1 *If* $\mathcal{C}_1 : A \Rightarrow B$ *and* $\mathcal{C}_2 : A \Rightarrow B$ *are channels, then so too are* $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)$ *and* $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$.

Proof This basically requires us to show that the bifunctions *left* and *right* associated with $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$ define homomorphisms, and similarly for $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)$. Let $(c_1 \| c_2) \in tok(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$ and $\gamma \in typ(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$ be such that $(c_1 \| c_2) : \gamma$ in $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2)$. Now, $\gamma = \langle \vee, \gamma_1, \gamma_2 \rangle$, for $\gamma_1 \in typ(\mathcal{C}_1)$ and $\gamma_2 \in typ(\mathcal{C}_2)$, and so $c_1 : \gamma_1$ in $\mathcal{C}_1$ or $c_2 : \gamma_2$ in $\mathcal{C}_2$. Suppose the former is the case; then, since $\mathcal{C}_1$ is a channel we can use the Principle of Harmony to show that it must be the case that $source(c_1) : ante(\gamma_1)$ and $target(c_1) : succ(\gamma_1)$ in $\mathcal{C}_1$. But $source(c_1 \| c_2) = source(c_1)$, so $source(c_1 \| c_2) : ante(\gamma_1)$. But, given the interaction between type-entailment and type-disjunction (Section 2.2.2), it must also be the case that $source(c_1 \| c_2) : (ante(\gamma_1) \vee ante(\gamma_2))$. Similarly, it can be shown that $target(c_1 \| c_2) : (succ(\gamma_1) \vee succ(\gamma_2))$. The case whereby $c_2 : \gamma_2$ in $\mathcal{C}_2$ leads to the same conclusion. Similarly, the analogous case for $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2)$ involving $\langle \wedge, \gamma_1, \gamma_2 \rangle$ leads to a similar property, which completes the proof of the desired result. □

The following result demonstrates associativity and commutativity of the parallel compositions. Associativity of serial composition is shown in (Seligman and Barwise 1993).

Proposition A.4.2 *For any channels* $\mathcal{C}_1$, $\mathcal{C}_2$ *and* $\mathcal{C}_3$ *linking classifications* A *and* B,

- $(\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) = (\mathcal{C}_2 \bigtriangleup \mathcal{C}_1)$, *and* $(\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) = (\mathcal{C}_2 \bigtriangledown \mathcal{C}_1)$;
- $(\mathcal{C}_1 \bigtriangleup (\mathcal{C}_2 \bigtriangleup \mathcal{C}_3)) = ((\mathcal{C}_1 \bigtriangleup \mathcal{C}_2) \bigtriangleup \mathcal{C}_3)$, *and* $(\mathcal{C}_1 \bigtriangledown (\mathcal{C}_2 \bigtriangledown \mathcal{C}_3)) = ((\mathcal{C}_1 \bigtriangledown \mathcal{C}_2) \bigtriangledown \mathcal{C}_3)$.

Proof The commutativity properties follow easily from the assumption that the type-conjunction and -disjunction operations on A and B are commutative. Associativity is also very easily checked (given associativity of the type-operations), although this is a lot more tedious. □

[8] The analogous result for serial composition is shown in (Seligman and Barwise 1993).

A.4.2 Relevant Logical Channels

The *Relevant Logical Channels* are based on *Relevant Logic* (e.g. (Anderson and Belnap 1975; Dunn 1984; Read 1988)) which is a particularly attractive non-classical logic. In particular, it shares with channel theory the underlying motivation of accounting for "information containment". Relevant implication is highly intensional and avoids the paradoxes of material implication: these paradoxes are the facts that any formula (materially) entails a tautology, and that a logical contradiction entails any formula. A Relevant Logical channel $\mathcal{R}$ effectively captures first-degree relevant entailment, in the following way: for token s and types ϕ, ψ, $(s : \phi)$ carries the information that $(s : \psi)$ via $\mathcal{R}$ iff ϕ' relevantly entails ψ', where ϕ', ψ' are logical sentences corresponding to the types ϕ, ψ.

Apart from the desirable properties associated with relevant implication I alluded to above, the logic of first-degree relevant entailment also has some interesting parallels with the operations of channel theory. Consider the following Hilbert-style axiomatisation of first-order entailment (Dunn 1984, p. 147):[9]

Axioms:

$A \wedge B \supset A, \quad A \wedge B \supset B$	Conjunction Elimination
$A \supset A \vee B, \quad B \supset A \vee B$	Disjunction Introduction
$A \wedge (B \vee C) \supset (A \wedge B) \vee C$	Distribution
$A \supset \neg\neg A, \quad \neg\neg A \supset A$	Double Negation

Rules of Inference:

$A \supset B, \quad B \supset C \quad \vdash \quad A \supset C$	Transitivity
$A \supset B, \quad A \supset C \quad \vdash \quad A \supset B \wedge C$	Conjunction Introduction
$A \supset C, \quad B \supset C \quad \vdash \quad A \vee B \supset C$	Disjunction Introduction
$A \supset B \quad \vdash \quad \neg B \supset \neg A$	Contraposition

The four rules of inference above correspond to the four Principles of Information flow that Barwise obtains from the operations of serial and parallel composition and contraposition (Barwise 1993), which in turn are obtained through the channel operations defined in Section 2.3.7. This means that we can define a Relevant Logical channel $\mathcal{R}$ simply by stipulating that there are constraints which correspond to the axioms,

[9]The following axiomatisation is for the first-degree fragment of Anderson and Belnap's system **R**.

and then closing $\mathcal{R}$ under the channel-theoretic operations, as required by the Logical Closure constraint.

Definition *Given a classification A, the Relevant Logical channel for A, denoted $\mathcal{R}_A : A \Rightarrow A$,*[10] *is the smallest logical channel whose set of types contains the following constraints, for each $\phi, \psi, \tau \in typ(A)$:*

- $(\phi \wedge \psi) \rightarrow \phi$
- $(\phi \wedge \psi) \rightarrow \psi$
- $\phi \rightarrow (\phi \vee \psi)$
- $\psi \rightarrow (\phi \vee \psi)$
- $(\phi \wedge (\psi \vee \tau)) \rightarrow ((\phi \wedge \psi) \vee \tau)$
- $\phi \rightarrow (\neg\neg\phi)$
- $(\neg\neg\phi) \rightarrow \phi$

Note that any Relevant Logical channel $\mathcal{R}$ must satisfy all the properties required of logical channels, as listed above. Hence, closure under the channel operations ensures that $\mathcal{R}$ models precisely the first-degree relevant entailments over a given set of types. This is shown by the following result.

Definition *Given a classification A with types $typ(A)$, a* propositional language based on *$typ(A)$ is a set $\mathcal{P}$ of propositional sentences such that, for each type $\phi \in typ(A)$, there is a sentence $\phi' \in \mathcal{P}$ satisfying the following constraints:*[11]

1. $(\neg\phi)' = \neg(\phi)'$;
2. $(\phi \wedge \psi)' = \phi' \wedge \psi'$;
3. $(\phi \vee \psi)' = \phi' \vee \psi'$.

Proposition A.4.3 *Let $\mathcal{R} : A \Rightarrow A$ be a relevant logical channel and L a propositional language based on A—given a type ϕ, ϕ' denotes the corresponding formula in L. Then $\phi \rightarrow \psi \in typ(\mathcal{R})$ iff $\phi' \supset \psi'$ is a valid first-degree entailment in* $\mathbf{R}$.

[10] Note that there is a different Relevant Logical channel for each classification A. I will usually drop the subscript on $\mathcal{R}_A$ when it causes no confusion.

[11] Some care should be taken here: the operations on the left-hand and right-hand sides of the equality sign are different of course—the one on the left is a type-operator while the one on the right is a logical symbol. This should cause no real problem.

Proof The proof is quite simple. I will simply show that the parallel composition operations preserve first-degree entailment—all other steps of the proof are even simpler.

Suppose $typ(\mathcal{R})$ contains $\phi_1 \rightarrow \psi_1$ and $\phi_2 \rightarrow \psi_2$, and that $\phi'_1 \supset \psi'_1$ and $\phi'_2 \supset \psi'_2$ are first-degree entailments. Using the parallel composition operations (under which $\mathcal{R}$ is closed), we obtain constraints $(\phi_1 \wedge \phi_2) \rightarrow (\psi_1 \wedge \psi_2)$, and $(\phi_1 \vee \phi_2) \rightarrow (\psi_1 \vee \psi_2)$. Consider the first of these (the other follows analogously)—we need to show that $(\phi'_1 \wedge \phi'_2) \supset (\psi'_1 \wedge \psi'_2)$ is a valid entailment in **R**. This is easily shown as follows: by Conjunction Elimination, we have $(\phi'_1 \wedge \phi'_2) \supset \phi'_1$ and $(\phi'_1 \wedge \phi'_2) \supset \phi'_2$; using Transitivity on each of these plus the two given entailments, we obtain $(\phi'_1 \wedge \phi'_2) \supset \psi'_1$ and $(\phi'_1 \wedge \phi'_2) \supset \psi'_2$, to which we apply Conjunction Introduction to obtain $(\phi'_1 \wedge \phi'_2) \supset (\psi'_1 \wedge \psi'_2)$, as required. $\square$

Appendix B

Channel Theoretic Model for Conditional Logics: Further Discussion and Results

This appendix contains extended treatments of some concepts which were covered only briefly in the main text of Chapter 3, as well as proofs of the formal results claimed in that chapter.

B.1 Nested Conditionals

In Section 3.1.2, I showed how conditional sentences can be interpreted within the channel theoretic analysis without any extension to the basic framework. However, the definition of the channel theoretic system for reasoning with conditionals basically ignored nested conditionals. In this section, I present one way in which the channel theoretic conditional logic can be extended so as to handle nested conditionals.

Jackson (1987) claims that the only meaningful nested conditionals are those that contain the inner conditional in the consequent of the outer one. Further, he claims that a nested conditional of the form "if X then if Y then Z" is *equivalent* to "if X and Y then Z". This observation underlies the following usual axiom of conditional logics:

$(A \rightarrow (B \rightarrow C)) \supset ((A \wedge B) \rightarrow C)$.

Basically, the antecedent of the outer conditional is effectively treated as an extra background condition of the inner conditional. This suggests the following treatment within the channel theoretic system for reasoning with conditionals: given a channel $\mathcal{C}$ supporting the outer conditional and a channel $\mathcal{C}'$ supporting the inner conditional, then there must be a channel $\mathcal{C}''$ in the hierarchy such that $\mathcal{C}' \sqsubseteq \mathcal{C}''$ and $\mathcal{C}''$ encodes the outer antecedent of $\mathcal{C}$ as a background assumption to $\mathcal{C}'$. For example, if $\mathcal{C}$ supports

> "If Pete doesn't have an umbrella then if it rains he'll get wet"

and $\mathcal{C}'$ supports "if it rains he'll get wet", then $\mathcal{C}''$ will effectively support the conditional

> "If Pete doesn't have an umbrella and it rains then he'll get wet."

The above behaviour can be imposed by defining the following constraint, along the lines of those defined in Section 3.5, which all channel hierarchies are required to satisfy.[1]

Definition Nested Conditional Constraint: *Let $\mathcal{C} : A \Rightarrow \mathcal{K}$ be a channel such that $\mathcal{K}$ is a classification of channels by conditional facts.*[2] *Suppose typ($\mathcal{C}$) contains the constraint $\phi {\rightarrow} \langle \Rightarrow, (b : \psi), (c : \tau) \rangle$ and tok($\mathcal{C}$) contains the connection $a \mapsto \mathcal{C}'$ where $\mathcal{C}' : B \Rightarrow C$. Then there is a channel $\mathcal{C}'' : (A \otimes B) \Rightarrow C$ such that $\mathcal{C}' \sqsubseteq_f \mathcal{C}''$ and $f(\psi {\rightarrow} \tau) = (\phi \wedge \psi) {\rightarrow} \tau$.*

This constraint ensures the axiom given earlier is supported by the channel theoretic logic of conditionals. The following result shows that Jackson's requirement that the two conditionals (one involving nesting, one conjoining the extra antecedent) are effectively equivalent holds.

Proposition B.1.1 *Suppose $(\phi {\rightarrow} \langle \Rightarrow, (b : \psi), (c : \tau) \rangle) \in$ typ($\mathcal{C}$) and $((\phi \wedge \psi) {\rightarrow} \tau) \in$ typ($\mathcal{C}''$). The information $(c : \tau)$ is carried in $\mathcal{C}$ via the former constraint iff $(c : \tau)$ is carried in $\mathcal{C}''$ via the latter constraint.*

Proof Suppose $(c : \tau)$ is carried via the constraint $\phi {\rightarrow} \langle \Rightarrow, (b : \psi), (c : \tau) \rangle$. Then $(a : \phi)$ holds in A and $(a \mapsto \mathcal{C}') : (\phi {\rightarrow} \langle \Rightarrow, (b : \psi), (c : \tau) \rangle)$ holds in $\mathcal{C}$. By the Principle of Harmony, this means that $(b \mapsto c : \psi {\rightarrow} \tau)$ must hold in $\mathcal{C}'$ which in turn means that $(b : \psi)$ must hold in B. But since $\mathcal{C}' \sqsubseteq \mathcal{C}'$ this must mean that $((\langle a, b \rangle \mapsto c) : (\phi \wedge \psi) {\rightarrow} \tau)$ must hold in $\mathcal{C}''$, so $(c : \tau)$ must hold in C—i.e. $(c : \tau)$ is carried via $(\phi \wedge \psi) {\rightarrow} \tau$ in $\mathcal{C}''$.

Conversely, suppose $(c : \tau)$ is carried via $(\phi \wedge \psi) {\rightarrow} \tau$ in $\mathcal{C}''$. From this it follows that

[1] The following definition is for the general setting—in particular, there is no assumption that channels are reflexive.

[2] I.e. the tokens of $\mathcal{K}$ are channels and the types are conditional facts.

$(a : \phi)$ holds in A and $(b : \psi)$ holds in B, from which the result follows. $\square$

B.2 Contextualising the Channel Theoretic Conditional Logic

The conditional channel operations of Section 3.4 take account of conflicting background conditions but only at the type level. The restriction of attention to the type-level is only reasonable under the assumption that channels and signalling relations are reflexive. For example, if $\mathcal{C} : A \Rightarrow B$, $\mathcal{C}' : B \Rightarrow C$, $\gamma \in typ(\mathcal{C})$ and $\gamma' \in typ(\mathcal{C}')$, then even if a background condition γ conflicts with a background condition of γ' *at the type level* this need not mean that the serial composition should in any way be affected since A and B are (presumably) dealing with different tokens. Of course, if the background condition on γ concerned the consequent of γ and the background condition of γ' concerned the antecedent of γ', then in this case type-conflict would need to be taken into account since the same classification B is involved in each case.

In the most general version of the channel theoretic framework—in which a channel may link two different classifications and signalling relations are not reflexive—it is *proposition* conflict/inconsistency rather than type-conflict that matters. Of course, this is true for the simplified case as well; it is just that under the simplification, proposition-inconsistency and type-conflict coincide. In general, however, proposition-conflict is a weaker notion than type-conflict—the *only* time they need coincide is when the same token is involved in both propositions. As such, the systems for reasoning with conditionals, generics and default properties of Chapters 3–5 can be seen as subsystems of the most general channel theoretic systems possible for these tasks.

Apart from the obvious simplification, the restriction to reflexive channels and signalling relations also allowed a more direct comparison to traditional "unsituated" logics of conditionals and default reasoning. In particular, I was able to show that the modified channel operations—which respect the background assumptions encoded in a hierarchy—satisfied many important patterns of inference (as well as invalidating various problematic ones). An investigation into the sorts of appropriate patterns of inference for a "fully contextualised" version of the channel theoretic systems (i.e. with the reflexivity simplification removed) would need to be undertaken in order to evaluate the resulting systems.

At this stage, the way in which the use of proposition conflict would replace the use of type-conflict in the conditional channel operations is unclear. A particular potential problem is that the background assumptions of a conditional may involve different tokens to those of the antecedent and consequent of the conditional. For example, consider the following modification to the example involving the British political parties.

> "If the Lib Dems get the most votes (in this election) then the Tories will be in power (for the next term)."
> "If the Tories aren't in power (for the next term) then Labour will be."

The second conditional is true only under the assumption that the Lib Dems don't get the most votes (in the current election). If we assume that the *current-election* situation and the *next-term* situation are different, then the superchannel encoding the background assumption of the second conditional involves tokens not necessarily available in the channel supporting this conditional. One way of overcoming this problem is to use complex classifications (and the corresponding subclassification relation) of the form described in Section A.2 and modify the definition of subchannel accordingly. For example, the channel $\mathcal{C} : T \Rightarrow T$ supporting the second conditional above would be a subchannel of the channel $\mathcal{C}' : (T \otimes E) \Rightarrow T$—$(T \otimes E)$ would contain tokens of the form $\langle t, e\rangle$, where t is a "term" situation and e an "election" situation, and $\mathcal{C}'$ would encode the background assumption in a constraint of the form $\langle \neg tories.in.power, \neg lib.dems.win\rangle \rightarrow labour.in.power$.

Other operations and conditions defined in Chapters 3–5 would also need to be modified. For example, the conditional Contraposition operation would require modification and the Consequent Consistency, Antecedent Background and Consistency Background Constraints would no longer be generally acceptable. The *maximal normality* condition of Chapter 5 would also need to be modified—specifically, some of the *inappropriateness* conditions would no longer reflect inappropriateness. Interestingly, the particular instances of inappropriateness that would need to be dropped (i.e. those requiring some form of consistency between antecedent and consequent types) are the same ones which I ignore in the description of how the methodology for defeasible reasoning relates to the maximal normality condition (Section 5.4.2), which reflects that that methodology is indeed a step towards *situated* defeasible reasoning.

B.3 The Problem with Necessitation Conditionals

Pollock (1976) has outlined a problem for the necessitation view of conditionals—whereby any conditional requires some sort of "link" between antecedent and consequent—which can be demonstrated by the following example.

> "If I press the button then the doorbell will ring."
> "If the doorbell rings then the doorbell exists."
>
> ---
>
> "If I press the button then the doorbell exists."

The first two conditionals are necessitation conditionals but the third is not—if it is true, it can only be because of the contingent truth of the fact that I have pressed the button and the bell exists. In particular, the third conditional does not seem to be a reasonable assertion under the channel theoretic analysis while the first two do.

The above pattern of inference is an instance of Transitivity that is validated using the conditional serial composition operation defined in Section 3.4 since the doorbell's non-existence is not an implicit assumption to the first conditional! It is the fact that the second conditional describes an entailment, and therefore has no associated background assumptions, that is at the root of the problem.[3] Barker (1994) has suggested that this problem involving entailment conditionals arises when the conclusion of the entailment conditional (i.e. the second one in the example), is an implicit assumption (though not a condition that necessarily holds) to the other conditional. For example, the assumption that the doorbell exists is implicit behind any assertion of the first conditional in the example.

Given this analysis of the problem, the operation of conditional serial composition as defined in Section 3.4.1 can be modified to avoid problematic examples of this sort by adding the extra proviso to the definition. (I have not explicitly defined what I mean be an "entailment channel": it may be taken to be that any channel satisfying the conditions required of logical channels—particularly that specified in Section 3.3.2—is an entailment channel.)

Definition *Let* $\mathcal{C}_1 : A \Rightarrow B$ *and* $\mathcal{C}_2 : B \Rightarrow C$ *be channels, with* $\gamma_1 \in typ(\mathcal{C}_1)$ *and* $\gamma_2 \in typ(\mathcal{C}_2)$*. The* conditional serial composition *of* $\mathcal{C}_1$ *and* $\mathcal{C}_2$*, written* $(\mathcal{C}_1 ; \mathcal{C}_2)$*, is identical*

[3]Note that any logic of conditionals with a possible-worlds semantics will also support the pattern of reasoning underlying this example.

to the standard serial composition of $\mathcal{C}_1$ and $\mathcal{C}_2$ except that $\langle\gamma_1, \gamma_2\rangle \in typ(\mathcal{C}_1\,;\,\mathcal{C}_2)$ only if

1. *there do not exist ${\mathcal{C}'}_1, {\mathcal{C}'}_2$ such that*
   1. *$\mathcal{C}_1 \sqsubseteq_{f_1} {\mathcal{C}'}_1$, $\mathcal{C}_2 \sqsubseteq_{f_2} {\mathcal{C}'}_2$; and*
   2. *$ante_{{\mathcal{C}'}_1}(f_1(\gamma_1)) \perp ante_{{\mathcal{C}'}_2}(f_2(\gamma_2))$; and*
3. *if $\mathcal{C}_2$ is an entailment channel, then there does not exist a channel ${\mathcal{C}'}_1$ such that $\mathcal{C}_1 \sqsubseteq_{f_1} {\mathcal{C}'}_1$ and $ante_{{\mathcal{C}'}_1}(f_1(\gamma_1))$ type-entails $succ_{\mathcal{C}_2}(\gamma_2)$.*

B.4 Proofs of Results in Chapter 3

In this section, I present proofs of all the propositions stated in Chapter 3. So as to simplify the presentation of the proofs, I will make use of the following notational convention. Let $\mathcal{C}$ and $\mathcal{C}'$ be channels, with $\phi{\rightarrow}\psi \in typ(\mathcal{C})$ and $\mathcal{C} \sqsubseteq_f \mathcal{C}'$. I will sometimes use $\phi^{f(\phi\rightarrow\psi)}$ to denote $ante(f(\phi{\rightarrow}\psi))$, which I will further abbreviate to ϕ^f when the constraint of which ϕ is the antecedent is unambiguously determined. Note, ϕ^f denotes the antecedent of the constraint to which f maps the constraint $\phi{\rightarrow}\psi$, of which ϕ is the antecedent. If there is more than one constraint, in the pertinent context, of which ϕ is an antecedent, then I will use the $\phi^{f(\phi\rightarrow\psi)}$ notation to disambiguate. Similarly, I will sometimes write ψ^f, and $\psi^{f(\phi\rightarrow\psi)}$, for $succ(f(\phi{\rightarrow}\psi))$.

The first result is from Section 3.4.3 and concerns the interaction of the conditional contraposition operation with the channel hierarchy.

Proposition 3.4.1 *Let $\mathcal{C} : A \Rightarrow B$ and $\mathcal{C}' : A' \Rightarrow B'$ be channels such that $\mathcal{C} \sqsubseteq_{f'} \mathcal{C}'$. Further suppose that $\phi{\rightarrow}\psi \in typ(\mathcal{C})$ and that $f(\phi{\rightarrow}\psi) = \phi'{\rightarrow}\psi'$.*

1. *If $\neg\psi{\rightarrow}\neg\phi \in typ(\mathcal{C}^*)$ then $\neg\psi'{\rightarrow}\neg\phi' \in typ(\mathcal{C}'^*)$;*
2. *$\mathcal{C}^* \sqsubseteq \mathcal{C}'^*$.*

Proof

1. Suppose $\neg\psi'{\rightarrow}\neg\phi' \notin typ(\mathcal{C}'^*)$. (For simplicity, I assume $\psi' = \psi$—the more general case is proved using an application of the Antecedent Background condition.) There exists a channel $\mathcal{C}''$ such that $\mathcal{C}' \sqsubseteq_{f'} \mathcal{C}''$, $f'(\phi'{\rightarrow}\psi') = \phi''{\rightarrow}\psi''$ and $\neg\psi' \perp \tau'$, where τ' is the least type such that $\phi'' \leq \phi' \wedge \tau'$. But the subchannel ordering is transitive, so we also have $\mathcal{C} \sqsubseteq_{f''} \mathcal{C}''$. Also, since $\phi \leq \phi'$ and $\psi = \psi'$, it is clearly

the case that $\neg\psi \perp \tau$, where τ is the least type such that $\phi'' \leq \phi \wedge \tau$. So the result follows.

2. This result follows pretty much from the previous one—all I basically need to show is that $(\neg\psi{\rightarrow}\neg\phi) \leq (\neg\psi'{\rightarrow}\neg\phi')$, and this is easily done.

□

The following results are from Section 3.5.3.

In a footnote to the definition of the function $(f_1 \parallel f_2)$ (Section 3.5.3), I mentioned that it was necessary to prove that the parallel composition of the higher channels (i.e. $(\mathcal{C}'_1 \parallel \mathcal{C}'_2)$) actually contains the type $\langle \wedge, f_1(\gamma_1), f_2(\gamma_2)\rangle$. This is shown as follows.

> **Proof** Let $f_1(\gamma_1)$ be γ_1' and let $f_2(\gamma_2)$ be γ_2', and suppose $\langle \wedge, \gamma_1', \gamma_2'\rangle$ is not contained in $(\mathcal{C}'_1 \parallel \mathcal{C}'_2)$. Then, by the definition of conditional parallel meet, there must be channels $\mathcal{C}''_1, \mathcal{C}''_2$ such that $\mathcal{C}'_1 \sqsubseteq_{f_1'} \mathcal{C}''_1$, $\mathcal{C}'_2 \sqsubseteq_{f_2'} \mathcal{C}''_2$, and $ante(f_1'(\gamma_1')) \perp ante(f_2'(\gamma_2'))$. But the subchannel relation is transitive, so we cannot then have $\langle \wedge, \gamma_1, \gamma_2\rangle$ contained in $(\mathcal{C}_1 \parallel \mathcal{C}_2)$, which contradicts the conditions stated in the definition. □

The next proposition regards the interaction of this function with the channel hierarchy.

Proposition 3.5.1 *Suppose all the conditions of the above definition hold. Then* $(\mathcal{C}_1 \parallel \mathcal{C}_2) \sqsubseteq_{(f_1 \parallel f_2)} (\mathcal{C}'_1 \parallel \mathcal{C}'_2)$.

Proof The proof involves checking that each condition of the subchannel definition holds. I will simply show that $\gamma \leq (f_1 \parallel f_2)(\gamma)$, for $\gamma \in typ(\mathcal{C} \parallel \mathcal{C}')$; the other conditions are checked more easily than this one.

Suppose γ is constructed from $\phi{\rightarrow}\psi \in typ(\mathcal{C}_1)$ and $\sigma{\rightarrow}\tau \in typ(\mathcal{C}_2)$. There are two possibilities:

1. Let γ be $(\phi \wedge \sigma){\rightarrow}(\psi \wedge \tau)$. By the definition of $\leq$, we have $\phi \leq \phi^{f_1}$ and $\psi^{f_1} \leq \psi$. Similarly, $\sigma \leq \sigma^{f_2}$ and $\tau^{f_2} \leq \tau$. So clearly, $(\phi \wedge \sigma) \leq (\phi^{f_1} \wedge \sigma^{f_2})$ and $(\psi^{f_1} \wedge \tau^{f_2}) \leq (\psi \wedge \tau)$; i.e. $(\phi \wedge \sigma){\rightarrow}(\psi \wedge \tau) \leq (\phi^{f_1} \wedge \sigma^{f_2}){\rightarrow}(\psi^{f_1} \wedge \tau^{f_2})$. But the latter is just $(f_1 \parallel f_2)(\gamma)$, so the result follows for this case;
2. Let γ be $(\phi \vee \sigma){\rightarrow}(\psi \vee \tau)$. This case follows exactly analogously to the other.

□

Proposition 3.5.2, which concerns the analogous result to the above for serial composition, is proved much the same way as the proposition above; the proof is even simpler.

The following result is from Section 3.5.5 and shows that channels failing the Reliability Constraint lack robustness.

Proposition 3.5.3 *Suppose in the definition that $ante(f_1(\phi \rightarrow \psi)) \perp ante(f_2(\phi \rightarrow \tau))$, and let $(a \mapsto b) \in tok(\mathcal{C})$ be such that $(a : \phi)$ holds in A. Then $a \mapsto b$ is an exception to at least one of the given constraints.*

Proof Let $c = a \mapsto b$ and suppose that $(c : \phi \rightarrow \psi)$ holds in $\mathcal{C}$. Then, by definition of the subchannel relation, $(c : f_1(\phi \rightarrow \psi))$ holds in $\mathcal{C}_1$. But then it cannot be the case that $(c : f_2(\phi \rightarrow \tau))$ holds in $\mathcal{C}_2$, due to the conflict in these types. Hence, again by definition of subchannel, it cannot be the case that $(c : \phi \rightarrow \tau)$ holds in $\mathcal{C}$—i.e. c is an exception. By reversing the roles of the constraints, we obtain the desired result. $\square$

The following result, from Section 3.5.6, concerns the lack of robustness of channels failing the Consequent Consistency Constraint.

Proposition 3.5.4 *Let $\mathcal{C}$ be a channel as described in the definition, and $\phi \rightarrow \psi \in typ(\mathcal{C})$ a constraint that fails the given condition. If $a \mapsto a \in tok(\mathcal{C})$ and $(a : \phi)$ holds in A, then $a \mapsto a$ is an exception to $\phi \rightarrow \psi$.*

Proof (The proposition involves the underlying assumption that a is coherent.) Let c be the connection $a \mapsto a$ and suppose that $(c : \phi \rightarrow \psi)$ holds. Then $(a : \psi)$ holds in A, which means $(a : \psi)$ holds in A' (since, by the definition of subchannel, $A \angle A'$). But we must also have that $(c : f(\phi \rightarrow \psi))$ holds, in which case $(a : ante(f(\phi \rightarrow \psi)))$ holds. But this can't be so without a being incoherent. $\square$

B.5 Evaluation of the Channel Theoretic Model—Proofs

This section contains the proofs that various axioms and rules of inference of standard conditional logics are supported by the Channel Theoretic model of conditionals (see Section 3.6.3).

Axioms

- *CC:* $((\phi \rightarrow \psi) \wedge (\phi \rightarrow \tau)) \supset (\phi \rightarrow (\psi \wedge \tau))$

  Proposition B.5.1 *Let $\mathcal{C}$ be a channel containing constraints $\phi \rightarrow \psi$ and $\phi \rightarrow \tau$. Then $(\mathcal{C} \parallel \mathcal{C})$ contains the constraint $\phi \rightarrow (\psi \wedge \tau)$.*

Proof $\phi \rightarrow (\psi \wedge \tau) \in typ(\mathcal{C} \parallel \mathcal{C})$ iff for every $\mathcal{C}_1, \mathcal{C}_2$ such that $\mathcal{C} \sqsubseteq_{f_1} \mathcal{C}_1$ and $\mathcal{C} \sqsubseteq_{f_2} \mathcal{C}_2$, it is not the case that $ante(f_1(\phi \rightarrow \psi)) \perp ante(f_2(\phi \rightarrow \tau))$. But this is exactly the (simpler version of the) Reliability condition, so the required result follows. □

- *RT:* $((\phi \rightarrow \psi) \wedge ((\phi \wedge \psi) \rightarrow \tau)) \supset (\phi \rightarrow \tau)$

Proposition B.5.2 *Let $\mathcal{C}$ be a channel containing constraints $\phi \rightarrow \psi$ and $(\phi \wedge \psi) \rightarrow \tau$. Then the channel $((\mathcal{C} \parallel \mathcal{L}) ; \mathcal{C})$ contains the constraint $\phi \rightarrow \tau$, where $\mathcal{L}$ is an appropriate logical channel.*[4]

Proof Consider $(\mathcal{C} \parallel \mathcal{L})$—by Reliability, $(\phi \rightarrow (\phi \wedge \psi)) \in typ(\mathcal{C} \parallel \mathcal{L})$.

To show that the composition of the constraints $\phi \rightarrow (\phi \wedge \psi)$ and $(\phi \wedge \psi) \rightarrow \tau$ is contained in $typ((\mathcal{C} \parallel \mathcal{L}) ; \mathcal{C})$, we need to ensure that, for any $\mathcal{C}_1, \mathcal{C}_2$ such that $(\mathcal{C} \parallel \mathcal{L}) \sqsubseteq_{f_1} \mathcal{C}_1$ and $\mathcal{C} \sqsubseteq_{f_2} \mathcal{C}_2$, it is not the case that $\phi^{f_1} \perp (\phi \wedge \psi)^{f_2}$, where $\phi^{f_1} = ante(f_1(\phi \rightarrow (\phi \wedge \psi)))$ and $(\phi \wedge \psi)^{f_2} = ante(f_2((\phi \wedge \psi) \rightarrow \tau))$. This is shown as follows.

By the Parallel Subchannel condition, $\mathcal{C}_1 = (\mathcal{C}' \parallel \mathcal{L}'')$, where $\mathcal{C} \sqsubseteq_{f'} \mathcal{C}'$ and $\mathcal{L} \sqsubseteq_{f''} \mathcal{L}''$. Let $\phi^{f'} = ante(f'(\phi \rightarrow \psi))$.[5] By the Reliability constraint, $\phi^{f'}$, ϕ and $(\phi \wedge \psi)^{f_2}$ must be mutually consistent. But $\phi^{f_1} = (\phi^{f'} \wedge \phi)$, since $f_1(\phi \rightarrow (\phi \wedge \psi))$ is formed by parallel composing $f'(\phi \rightarrow \psi)$ and $\phi \rightarrow \phi$. Hence, it can't be the case that $\phi_1^{f_1}$ is inconsistent with $(\phi \wedge \psi)^{f_2}$. This completes the proof of the result. □

The above is perhaps best visualised by way of an illustration; see Figure B.1.

- *ST10:* $(\phi \rightarrow \psi) \supset (\neg\neg\phi \rightarrow \psi)$

Proposition B.5.3 *Let $\mathcal{C}$ be a channel containing the constraint $\phi \rightarrow \psi$. Then the channel $(\mathcal{L} ; \mathcal{C})$ contains the constraint $\neg\neg\phi \rightarrow \psi$, where $\mathcal{L}$ is an appropriate logical channel.*

Proof Suppose there is a $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$ and $ante(f(\phi \rightarrow \psi)) \perp \neg\neg\phi$, where $\rightarrow\neg\neg\phi \in typ(\mathcal{L})$. By the minimal conditions on the $\perp$ relation (Section 2.2.2), it must be the case that $\phi \perp ante(f(\phi \rightarrow \psi))$ (since $\phi \leq \neg\neg\phi$), which contradicts the Reliability Constraint. So the required result follows. □

[4]I.e. $\mathcal{L}$ is the logical channel linking the same classifications as does $\mathcal{C}$.

[5]Since $\mathcal{L}$ is a logical channel, it must (by definition) be the case that $f''(\phi \rightarrow \phi) = \phi \rightarrow \phi$.

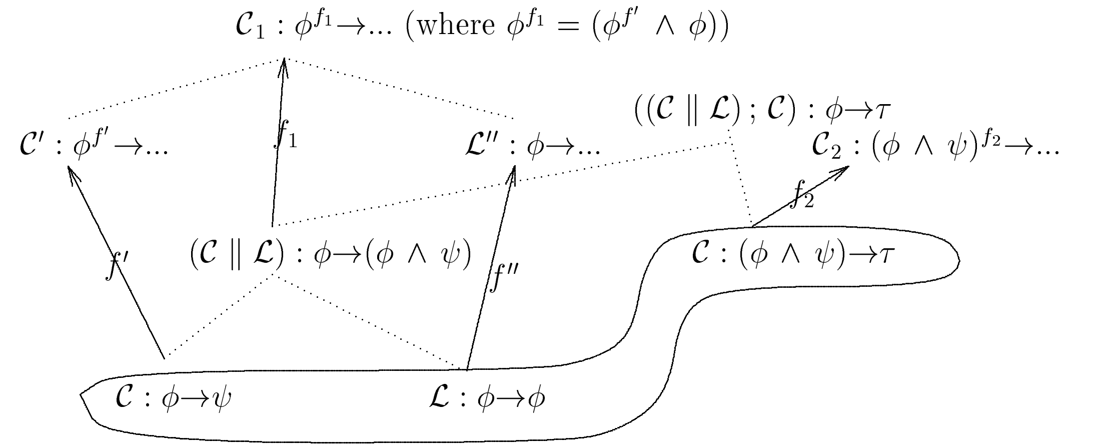

Figure B.1: Companion to proof ...

Rules of inference

- *RCEC: if* $\phi \iff \psi$ *then* $(\phi\!\rightarrow\!\tau) \iff (\psi\!\rightarrow\!\tau)$

  Proposition B.5.4 *Let* $\mathcal{L} : A \Rightarrow A$ *be a logical channel containing* $\phi\!\rightarrow\!\psi$ *and* $\psi\!\rightarrow\!\phi$. *If* $\mathcal{C} : A \Rightarrow A$ *contains* $\phi\!\rightarrow\!\tau$ *(resp.,* $\psi\!\rightarrow\!\tau$*) then* $(\mathcal{L}\,;\,\mathcal{C})$ *contains* $\psi\!\rightarrow\!\tau$ *(resp.,* $\phi\!\rightarrow\!\tau$*).*

  Proof The result holds so long as there is no channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$ and $\psi \perp ante(f(\phi\!\rightarrow\!\tau))$. This must be the case as otherwise the Reliability constraint would be violated, since $\phi \leq_A \psi$. (The other case is symmetric.) $\square$

- *RCEA: if* $\phi \iff \psi$ *then* $(\tau\!\rightarrow\!\phi) \iff (\tau\!\rightarrow\!\psi)$

  Proposition B.5.5 *Let* $\mathcal{L} : A \Rightarrow A$ *be a logical channel containing* $\phi\!\rightarrow\!\psi$ *and* $\psi\!\rightarrow\!\phi$. *If* $\mathcal{C} : A \Rightarrow A$ *contains* $\tau\!\rightarrow\!\phi$ *(resp.,* $\tau\!\rightarrow\!\psi$*) then* $(\mathcal{C}\,;\,\mathcal{L})$ *contains* $\tau\!\rightarrow\!\psi$ *(resp.,* $\tau\!\rightarrow\!\phi$*).*

  Proof The result holds so long as there is no channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$ and $\phi \perp ante(f(\tau\!\rightarrow\!\phi))$. This must be the case as otherwise the Consequent Consistency constraint would be violated. $\square$

- *RCK: if* $(\phi_1 \wedge ... \wedge \phi_n) \supset \psi$ *then* $((\tau\!\rightarrow\!\phi_1) \wedge ... \wedge (\tau\!\rightarrow\!\phi_n)) \supset (\tau\!\rightarrow\!\psi)$

Proposition B.5.6 *Let* $\mathcal{L} : A \to A$ *be a logical channel containing* $(\phi_1 \wedge ... \wedge \phi_n) \to \psi$. *If* $\mathcal{C} : A \to A$ *is a channel containing* $\tau \to \phi_1 \ ... \ \tau \to \phi_n$ *then* $((\mathcal{C} \parallel ... \parallel \mathcal{C}) ; \mathcal{L})$ *contains* $\tau \to \psi$.

Proof By Reliability, it is clear that $(\mathcal{C} \parallel ... \parallel \mathcal{C})$ must contain $\tau \to (\phi_1 \wedge ... \wedge \phi_n)$. Let $\mathcal{C}'$ be a channel such that $(\mathcal{C} \parallel ... \parallel \mathcal{C}) \sqsubseteq_f \mathcal{C}'$. By the Consequent Consistency constraint, it cannot be the case that $ante(f(\tau \to (\phi_1 \wedge ... \wedge \phi_n))) \perp (\phi_1 \wedge ... \wedge \phi_n)$, and so the result follows. $\square$

Appendix C

A Channel Theoretic Model for Generics: Further Discussion and Results

This appendix contains an illustration of the use of Morreau's logic of commonsense entailment.

C.1 An Illustration of Commonsense Entailment

In Section 4.1.4, I noted that Morreau's logic of commonsense entailment does not fall victim to the problem of Irrelevance. For example, $\forall x(bird(x) \wedge red(x) \rightarrow flies(x))$ can be (defeasibly) inferred from $\forall x(bird(x) \rightarrow flies(x))$. This is illustrated here. For simplicity, I will consider only the propositional variant of this example.

To show that $(bird \rightarrow flies) \approx\!\!\!\!\mid\; (bird \wedge red \rightarrow flies)$, we need to show that for every fixpoint t accessible (via some chain of normalisations) from $\Theta + \{(bird \rightarrow flies)\}$, $t \models (red \wedge bird \rightarrow flies)$.[1] Let s be the state $\Theta + \{(bird \rightarrow flies)\}$. This state contains exactly those worlds w such that $^*(w, \|bird\|) \subseteq \|flies\|$. The normalisation of s with respect to $\|bird\|$ is defined as follows.

[1]Even though I do not explicitly state it, recall that the collection of worlds, individuals, the accessibility function * and the interpretation $\|.\|$ of non-logical symbols are from the canonical model.

It is easily shown that (in the canonical model) $^*(s, \|bird\|) \cap s \neq \emptyset$, so $\mathcal{N}(s, \|bird\|) = s \backslash (\|bird\| \backslash ^*(s, \|bird\|))$. Now, $\|bird\| \backslash ^*(s, \|bird\|)$ is the collection of worlds that support $\|bird\|$ yet are not normal *bird* worlds. So removing such worlds from s leaves us with a collection of worlds W such that each world in W supports $bird \rightarrow flies$ (all worlds in s support this sentence already) and any world in W that supports *bird* also supports any property that normal birds possess—i.e. for every $w \in W$ such that $w \in \|bird\|$, $w \in {}^*(s, \|bird\|)$; in particular, if $w \in \|bird\|$, then $w \in \|flies\|$, since (from above) $^*(s, \|bird\|) \subseteq \|flies\|$. It turns out that s is actually a fixpoint of normalisation and $\|bird\|$ is the only proposition with respect to which we need to normalise t, so s is the only information state we need to consider.

We now need to determine whether s (defeasibly) supports the sentence $(bird \wedge red \rightarrow flies)$. This is true iff $w \models (bird \wedge red \rightarrow flies)$ for every $w \in s$, which in turn is true iff $^*(w, \|bird \wedge red\|) \subseteq \|flies\|$ for every $w \in s$. By *Facticity,* we know that $^*(w, \|bird \wedge red\|) \subseteq \|bird \wedge red\|$;[2] hence $^*(w, \|bird \wedge red\|) \subseteq \|bird\|$. But for every $w \in s$, we have removed worlds for which $w \in \|bird\|$ and $w \notin {}^*(s, \|bird\|)$. So, if $w \in s$ and $w \in \|bird\|$, then $w \in {}^*(s, \|bird\|)$, and so $w \in \|flies\|$. So it follows that, for all $w \in s$, $^*(w, \|bird \wedge red\|) \subseteq \|flies\|$; i.e. $s \models (bird \wedge red \rightarrow flies)$. This proves that

$$(bird \rightarrow flies) \approx\!\!\!\!| \; (bird \wedge red \rightarrow flies).$$

[2] All "normal" red birds are red birds!

Appendix D

A Channel Theoretic Model of Default Reasoning: Further Discussion and Results

This appendix contains proofs of the results claimed in Chapter 5.

D.1 Proofs of Propositions from Chapter 5

This section contains the proofs of several results from Chapter 5.

Proposition 5.2.1 *For any channels $\mathcal{C}$ and $\mathcal{C}'$, if $\mathcal{C} \sqsubseteq_f \mathcal{C}'$ then $Link(\mathcal{C}) \preceq_f Link(\mathcal{C}')$.*

Proof We simply need to check each of the conditions in the definition of sublink.

1. The first condition clearly holds, by definition of subchannel;
2. Suppose $\langle t, t' \rangle \in Link(\mathcal{C})^{\vee}$. Then there exists $c \in tok(\mathcal{C})$ such that $source_{\mathcal{C}}(c) = t, target_{\mathcal{C}}(c) = t'$. But, by definition of subchannel, $c \in tok(\mathcal{C}')$ and $source_{\mathcal{C}'}(c) = source_{\mathcal{C}}(c), target_{\mathcal{C}'}(c) = target_{\mathcal{C}}(c)$. Hence, the second condition follows;
3. Suppose $\langle \phi, \psi \rangle \in Link(\mathcal{C})^{\wedge}$. Then there exists $\gamma \in typ(\mathcal{C})$ such that $ante(\gamma) = \phi, succ(\gamma) = \psi$. By definition of subchannel, $\gamma \in typ(\mathcal{C}')$, $ante_{\mathcal{C}}(\gamma) \leq_{A'} ante_{\mathcal{C}'}(\gamma)$ and $succ_{\mathcal{C}'}(\gamma) \leq_{B'} succ_{\mathcal{C}}(\gamma)$. Hence, the third condition follows.

□

Proposition 5.2.2 *Let $\mathcal{D} = \langle \mathcal{L}, \mathcal{F} \rangle$ be a link hierarchy and let $\mathcal{D}' = \langle \mathcal{L}', \mathcal{F} \rangle$ be a maximally normal version of $\mathcal{D}$. Further, let $L_1, L_2 \in \mathcal{L}$, where $L_1 \preceq_f L_2$, and let $L_1', L_2' \in \mathcal{L}'$ be the links corresponding to L_1, L_2 respectively. Then $L_1' \preceq_f L_2'$.*

Proof Since L_1', L_2' are exactly the same as L_1, L_2 except possibly for the signalling relations, then I only need to show that $(L_1')^{\vee} \subseteq (L_2')^{\vee}$. Suppose $\langle a, b \rangle \notin L_2'$. Then $\langle a, b \rangle$ is internally inappropriate in L_2'. So one of the following cases must hold:

1. $(b : \psi)$ holds, such that $\langle \phi_2, \psi_2 \rangle \in (L_2')^{\wedge}$ and $\psi \perp \psi_2$. In this case, there is some $\langle \phi_1, \psi_1 \rangle \in {L_1}^{\vee}$, such that $f(\langle \phi_1, \psi_1 \rangle) = \langle \phi_2, \psi_2 \rangle$.[1] Hence, $\psi_2 \leq \psi_1$ and so $\psi \perp \psi_1$, which means that $\langle a, b \rangle$ is inappropriate in L_1', and so $\langle a, b \rangle \notin (L_1')^{\vee}$.
2. $(a : \phi)$ holds, such that $\langle \phi_2, \psi_2 \rangle \in (L_2'')^{\wedge}$, where $L_2' \preceq L_2''$ and $\phi \perp \phi_2$. By the transitivity of the subchannel and sublink relations, the result follows easily in this case.

So in either case, $\langle a, b \rangle$ is inappropriate in L_1', and the required result follows. □

D.2 Channel-Theoretic Default Systems and Preference Relations—Proofs

This section contains the proof of Theorem 5.3.2 in Section 5.3.3, showing that channel-theoretic default consequence corresponds to preferential consequence. The following result proves useful in the proof of the theorem.

Lemma *Let L and L' be links such that $\langle \alpha, \beta \rangle \in L^{\wedge}$, $L \preceq_f L'$ and $f(\langle \alpha, \beta \rangle) = \langle \alpha', \beta' \rangle$. Further suppose that $(a : \alpha)$ holds and that $\langle a, a \rangle$ is not internally inappropriate in L.*

1. *If $\alpha' \perp \alpha$ then $\langle a, a \rangle$ is externally inappropriate in L.*
2. *If $\alpha' \perp \beta$ then $\langle a, a \rangle$ is externally inappropriate in L.*

Proof The proof is straightforward for each case, using L as both links required in the definition of external inappropriateness, and using the fact that $L \preceq L$ for the first case. □

Theorem 5.3.2 *For any collection $\mathcal{S}$ of channels, $\mid\!\sim_{\mathcal{S}}$ satisfies all the conditions corresponding to the principles required of preferential consequence relations.*

[1]This holds for default hierarchies since each (non-logical) channel of a default hierarchy contains exactly one constraint.

Proof To prove the theorem, I need to consider in turn the channel-theoretic version of each of *KLM*'s constraints and show that each one is satisfied by the $\mid\!\sim_{\mathcal{S}}$ consequence relation.

Recall that the Parallel and Serial Subchannel Constraints ensure that if $\mathcal{C}$ is a channel formed by composition from $\mathcal{C}_1$ and $\mathcal{C}_2$, and $\mathcal{C} \sqsubseteq \mathcal{C}'$, then $\mathcal{C}'$ is itself composed from channels $\mathcal{C}'_1$ and $\mathcal{C}'_2$, where $\mathcal{C}_1 \sqsubseteq \mathcal{C}'_1$ and $\mathcal{C}_2 \sqsubseteq \mathcal{C}'_2$. In particular, this means that the background assumptions of $\mathcal{C}'$ are related to the background assumptions of $\mathcal{C}$. For example, if $\mathcal{C}$ is formed from parallel composition, then the antecedent of any constraint in $\mathcal{C}'$ must be of the form $\alpha' \wedge \beta'$, where α' (resp. β') is the antecedent of a constraint from $\mathcal{C}'_1$ (resp. $\mathcal{C}'_2$). Similarly, since a logical channel $\mathcal{L}$ has no superchannels (other than itself), then the background assumptions of the channel formed from the composition of a logical channel with some other channel $\mathcal{C}$ must be the background assumptions of $\mathcal{C}$. For example, if $\mathcal{L}$ contains the constraint $\alpha \rightarrow \beta$, $\mathcal{C}$ contains the constraint $\tau \rightarrow \sigma$, and $(\mathcal{L} \,\|\, \mathcal{C}) \sqsubseteq \mathcal{X}$, then the antecedent of any constraint in $\mathcal{X}$ must be of the form $\alpha \wedge \tau'$, where τ' is the antecedent of a constraint in some superchannel of $\mathcal{C}$.

Left Logical Equivalence. Given that the logical link $Link(\mathcal{L})$ contains both $\langle \alpha, \beta \rangle$ and $\langle \beta, \alpha \rangle$ in its signalling relation, and that $(a : \alpha) \mid\!\sim (a : \gamma)$, I need to show that $(a : \beta) \mid\!\sim (a : \gamma)$. This holds if and only if there is a link in the default hierarchy whose signalling relation contains $\langle a, a \rangle$ and whose indicating relation contains $\langle \beta, \gamma \rangle$. Given the premises, there must be a link L such that $L^{\wedge}$ contains $\langle \alpha, \gamma \rangle$.

Let $\mathcal{C}$ be the channel from which L is constructed (i.e. $L = Link(\mathcal{C})$) and suppose that $(\mathcal{L}\,;\,\mathcal{C})$ does not contain $\langle \beta, \gamma \rangle$. This can only be the case if there is a channel $\mathcal{C}'$ such that $\mathcal{C} \sqsubseteq_f \mathcal{C}'$, where $f(\alpha \rightarrow \gamma) = \alpha' \rightarrow \gamma'$ and $\alpha' \perp \beta$. But (from the constraints supported in the logical channel) we know that $\beta \leq \alpha$, so we then must have $\alpha \perp \alpha'$, which contradicts the result of the lemma above since $\langle a, a \rangle$ cannot be externally inappropriate in L. So the indicating relation of $Link(\mathcal{L}\,;\,\mathcal{C})$ (call this link LC) must contain $\langle \beta, \gamma \rangle$.

To show that $\langle a, a \rangle \in LC^{\vee}$, I need to show that $\langle a, a \rangle$ is not inappropriate in LC. This is done by considering each of the cases (one each for internal and external inappropriateness) in turn.

1. Suppose $\langle a, a \rangle$ is internally inappropriate in LC. Then either (i) $(a : \phi)$ holds and $\phi \perp \gamma$, in which case $\langle a, a \rangle$ must also be internally inappropriate in L, contradicting $(a : \alpha) \mid\!\sim (a : \gamma)$; or (ii) there exists a link LC' such that $LC \preceq_f LC'$, $f(\langle \beta, \gamma \rangle) = \langle \beta', \gamma' \rangle$, $(a : \beta'')$ holds and $\beta'' \perp \beta'$. By the nature of the subchannel (and therefore

sublink) relation's behaviour w.r.t. sequential composition, it must be the case that there is a link L' such that $L \preceq_{f'} L'$, $f'(\langle\alpha,\gamma\rangle) = \langle\alpha',\gamma'\rangle$ and $\beta' \leq \beta \wedge \alpha'$ (since the logical channel/link has no associated background conditions). (This is illustrated in Figure D.1.) But $\beta \leq \alpha$ and $\alpha \leq \alpha'$, so $\beta \wedge \alpha' = \alpha'$. Hence, if $\beta' \perp \beta''$ then $\alpha' \perp \beta''$. But then $\langle a,a\rangle$ would be internally inappropriate in L, which cannot be since L supports the inference $(a : \alpha) \mid\!\!\sim\!\!\!> (a : \gamma)$ and hence $L^{\vee}$ must contain $\langle a,a\rangle$. So $\langle a,a\rangle$ cannot be internally inappropriate in LC.

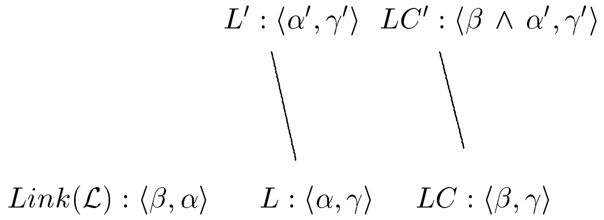

Figure D.1: Channel hierarchy when $\langle a,a\rangle$ is internally inappropriate in LC

2. Suppose $\langle a,a\rangle$ is externally inappropriate in LC. Then there must be links LC', N, N'[2] such that $LC \preceq_f LC'$, $N \preceq_{f'} N'$, $\langle\tau,\sigma\rangle \in N^{\wedge}$—as illustrated by Figure D.2—and

1. $\langle a,a\rangle$ is not internally inappropriate in N;

2. $f(\langle\beta,\gamma\rangle) = \langle\beta',\gamma'\rangle$ and $f'(\langle\tau,\sigma\rangle) = \langle\tau',\sigma'\rangle$;

3. $(t : \beta)$ holds and either $(t : \tau)$ also holds or $\tau \leq \gamma$; and

4. either $\beta' \perp \sigma$, $\tau' \perp \gamma$ or $\beta' \perp \tau'$.

(I will only show the result for the case where $(t : \beta)$ and $(t : \tau)$ both hold. The proof for the case where $(t : \beta)$ holds and $\tau \leq \gamma$ follows more easily.) Suppose the first three items hold and consider the conditions described in the last. Since $\langle a,a\rangle \in L^{\wedge}$, $\langle a,a\rangle$ cannot be externally inappropriate in L, it cannot be the case that $\tau' \perp \gamma$. Similarly, as described in the proof regarding internal inappropriateness, any background conditions to LC must come from some superlink L' of L satisfying the conditions described above. Again, $\langle a,a\rangle$ is not externally inappropriate in L so it cannot be the case that $\alpha' \perp \sigma$, and since I showed that $\beta' \leq \alpha'$ it also cannot be the case that $\beta' \perp \sigma$. Similarly, it cannot be the case that $\alpha' \perp \tau'$, so

[2]Actually, to be more precise I need to show that there cannot be n channels satisfying the requirements of the more complex definition of external inappropriateness. However, this greatly complicates the presentation and the proof generalises easily enough.

$$\begin{tikzcd} LC' : \langle \beta', \gamma' \rangle \arrow[d, dash, "f"] & N' : \langle \tau', \sigma' \rangle \arrow[d, dash, "f'"] \\ LC : \langle \beta, \gamma \rangle & N : \langle \tau, \sigma \rangle \end{tikzcd}$$

$$\beta' \perp \sigma \text{ or } \tau' \perp \gamma \text{ or } \beta' \perp \tau'$$

Figure D.2: Conditions for $\langle a, a \rangle$ to be externally inappropriate in LC

neither can it be the case that $\beta' \perp \tau'$. This covers all the possible cases associated with external inappropriateness, so $\langle a, a \rangle$ cannot be externally inappropriate in LC.

So, I have shown that $\langle \beta, \gamma \rangle \in LC^{\wedge}$ and that $\langle a, a \rangle$ is not inappropriate in LC, in which case (by maximal normality) $\langle a, a \rangle \in LC^{\vee}$. Hence, $(a : \beta) \mid\!\!\rightsquigarrow (a : \gamma)$, which proves that $\mid\!\!\rightsquigarrow$ satisfies Left Logical Equivalence.

The rest of the proof is not described in such detail as the above case.

Right Weakening. Suppose $L = Link(\mathcal{C})$ is the link supporting $(a : \gamma) \mid\!\!\rightsquigarrow (a : \alpha)$, and let CL be the link $Link(\mathcal{C} \,;\, \mathcal{L})$, where $\mathcal{L}$ is the logical channel (and $\langle \alpha, \beta \rangle \in Link(\mathcal{L})^{\wedge}$). I will show that $\langle \gamma, \beta \rangle \in CL^{\wedge}$ and $\langle a, a \rangle \in CL^{\vee}$.

Suppose that $\langle \gamma, \beta \rangle \notin CL^{\wedge}$. Then there must be a link L' such that $L \preceq_f L'$, $f(\langle \gamma, \alpha \rangle) = \langle \gamma', \alpha' \rangle$ and $\gamma' \perp \alpha$. But then $\langle a, a \rangle$ would be externally inappropriate in L (using the result of the above lemma), which contradicts $\langle a, a \rangle \in L^{\wedge}$.

Suppose that $\langle a, a \rangle$ is internally inappropriate in CL. Then there is a link CL' such that $CL \preceq_f CL'$, $f(\langle \gamma, \beta \rangle) = \langle \gamma', \beta' \rangle$ and either (i) $(a : \phi)$ holds and $\phi \perp \beta$, or (ii) $(a : \phi)$ holds and $\phi \perp \gamma'$. If case (i) holds, then since $\beta \leq \alpha$, then $\phi \perp \alpha$, in which case $\langle a, a \rangle$ would be internally inappropriate in L, contradicting $\langle a, a \rangle \in L^{\vee}$. If case (ii) holds, then again $\langle a, a \rangle$ would be internally inappropriate in L, so $\langle a, a \rangle$ cannot be internally inappropriate in CL.

Suppose that $\langle a, a \rangle$ is externally inappropriate in CL. Then there is a link with the properties of CL' above, and links N and N' with the properties described earlier (i.e. in the proof for LLE) such that (among the other requirements of the definition of external inappropriateness), either $\gamma' \perp \sigma$ or $\tau' \perp \beta$ or $\tau' \perp \gamma'$. Since $\langle a, a \rangle$ is not externally inappropriate in L, it can neither be the case that $\tau' \perp \alpha$ nor that $\gamma' \perp \sigma$ nor that $\gamma' \perp \tau'$ (recall that the logical link does not have any background assumptions). Since $\beta \leq \alpha$, neither can it be the case that $\tau' \perp \beta$. Hence, all possible cases of external inappropriateness are covered, completing the proof for the case of Right Weakening.

Reflexivity: This constraint is trivially supported since the logical channel contains the constraint $\alpha \rightarrow \alpha$ for every type α.

And: Let $L_1 = Link(\mathcal{C}_1)$ and $L_2 = Link(\mathcal{C}_2)$ be the links that support $(a : \alpha) \mathrel{|\!\!\leadsto} (a : \beta)$ and $(a : \alpha) \mathrel{|\!\!\leadsto} (a : \gamma)$ respectively. I will show that $L = Link(\mathcal{C}_1 \parallel \mathcal{C}_2)$ supports $(a : \alpha) \mathrel{|\!\!\leadsto} (a : \beta \wedge \gamma)$.

Suppose that $\langle \alpha, (\beta \wedge \gamma) \rangle \notin L^{\wedge}$. Then there are links L_1', L_2' such that $L_1 \preceq_{f_1} L_1'$, $L_2 \preceq_{f_2} L_2'$, $f_1(\langle \alpha, \beta \rangle) = \langle \alpha', \beta' \rangle$, $f_2(\langle \alpha, \gamma \rangle) = \langle \alpha'', \gamma' \rangle$ and $\alpha' \perp \alpha''$. But this cannot be the case as otherwise $\langle a, a \rangle$ would be externally inappropriate in these links, unless it was internally inappropriate in one of them—either way, this contradicts the fact that $\langle a, a \rangle \in L_1^{\vee}, L_2^{\vee}$.

Suppose $\langle a, a \rangle$ is internally inappropriate in L. Then $(a : \phi)$ holds and either (i) $\phi \perp (\beta \wedge \gamma)$ or (ii) there is a link L' such that $L \preceq_f L'$, $f(\langle \alpha, \beta \wedge \gamma \rangle) = \langle \alpha', \psi \rangle$ and $\phi \perp \alpha'$. Since all that is assumed in the initial classification is $(a : \alpha)$, then for $(a : \phi)$ to hold it must be the case that $\phi \leq \alpha$, in which case it must also be that $\alpha' \perp \alpha$. But this contradicts the above lemma, so (ii) cannot hold. Further, if $\phi \perp (\beta \wedge \gamma)$ then $\alpha \perp (\beta \wedge \gamma)$, in which case $\langle a, a \rangle$ is externally inappropriate in each of L_1 and L_2 (using the complex definition of external inappropriateness), so (i) also leads to a contradiction.

Suppose $\langle a, a \rangle$ is externally inappropriate in L. Then there is a link L' with the properties above and the usual links N and N' described earlier, such that $\alpha' \perp \sigma$, $\tau' \perp (\beta \wedge \gamma)$ or $\alpha' \perp \tau'$. From what was said in the paragraph preceding the proof of *LLE,* it must be the case that $\alpha' = (\alpha_1' \wedge \alpha_2')$, where α_1', α_2' are antecedents of constraints in super-links of L_1 and L_2 respectively. Hence, by the (complex) definition of external inappropriateness, it can neither be the case that $\alpha' \perp \sigma$ nor that $\alpha' \perp \tau'$, since otherwise $\langle a, a \rangle$ would be externally inappropriate in L_1 and L_2 (and N). Similarly, if it was the case that $(\beta \wedge \gamma) \perp \tau'$ then once again $\langle a, a \rangle$ would be externally inappropriate in L_1 and L_2. Hence $\langle a, a \rangle$ cannot be externally inappropriate in L.

Or: Let L_1, L_2 and L be as in the proof for *And,* with the difference that L_1 supports $(a : \alpha) \mathrel{|\!\!\leadsto} (a : \gamma)$ and L_2 supports $(a : \beta) \mathrel{|\!\!\leadsto} (a : \gamma)$. By the nature of conditional parallel composition, it is clear that $\langle \alpha \vee \beta, \gamma \rangle \in L^{\wedge}$.

It is also easily shown that $\langle a, a \rangle$ is neither internally nor externally inappropriate in L. By the nature of the interaction of conditional parallel composition with the subchannel relation, the antecedent of any constraint in a superchannel of L must be of the form $\alpha' \vee \beta'$, where α' and β' are antecedents of constraints in superchannels of L_1 and L_2 respectively. Since $(\alpha' \vee \beta') \leq \alpha'$ (resp., β'), then it is easily checked that for $\langle a, a \rangle$ to be internally/externally inappropriate in L it must also be internally/externally inappropriate in L_1 (resp., L_2), which leads to a contradiction.

Cautious Monotonicity: Let $L_1 = Link(\mathcal{C}_1)$ be the link supporting $(a : \alpha) \mathrel{|\!\!\sim\!\!\!>} (a : \gamma)$, $L_2 = Link(\mathcal{C}_2)$ be the link supporting $(a : \alpha) \mathrel{|\!\!\sim\!\!\!>} (a : \beta)$ and $\mathcal{L}$ be the logical channel (containing the constraint $\alpha \wedge \beta \rightarrow \alpha$). I will show that the link $LC = Link(\mathcal{L}\,;\mathcal{C}_1)$ supports the required inference.

Suppose LC does not contain $\langle \alpha \wedge \beta, \gamma \rangle$. Then there must be a link L_1' such that $L_1 \preceq_f L_1'$, $f(\langle \alpha, \gamma \rangle) = \langle \alpha', \gamma' \rangle$ and $\alpha' \perp (\alpha \wedge \beta)$. However, this would mean that $\langle a, a \rangle$ was externally inappropriate in L_1, which cannot be the case.

Suppose that $\langle a, a \rangle$ is internally inappropriate in LC. Then there is a link LC' such that $LC \preceq_f LC'$, $f(\langle \alpha \wedge \beta, \gamma \rangle) = \langle \alpha' \wedge \beta, \gamma' \rangle$ (given the interaction of the subchannel hierarchy with serial composition and the fact that the logical channel has no background conditions), and either (i) $(a : \phi)$ and $\phi \perp (\alpha' \wedge \beta)$; or (ii) $(a : \phi)$ and $\phi \perp \gamma$. Clearly, (ii) cannot be the case as $\langle a, a \rangle$ would otherwise be internally inappropriate in L_1. Also, since $(a : \alpha \wedge \beta)$ is the only premise then it must be that $\phi \leq (\alpha \wedge \beta)$, hence (i) would require that $(\alpha \wedge \beta) \perp (\alpha' \wedge \beta)$. This would mean that $\langle a, a \rangle$ was externally inappropriate in $Link(\mathcal{L})$, L_1 and L_2, which cannot be the case.

Suppose that $\langle a, a \rangle$ is externally inappropriate in LC. Then there is a link LC' as above and the usual links N, N' such that either $(\alpha' \wedge \beta) \perp \sigma$ or $\tau' \perp \gamma$ or $(\alpha' \wedge \beta) \perp \tau'$. As usual, if any of these was the case, then $\langle a, a \rangle$ would also be externally inappropriate in one or both of L_1 and L_2, which leads to a contradiction. The only case that is non-trivial is that involving $\tau' \perp \gamma$. Since $(a : \tau)$ must hold (if $\langle a, a \rangle$ is externally inappropriate in N) and $(a : \alpha \wedge \beta)$ is the only premised information, then it must be the case that $\tau \leq (\alpha \wedge \beta)$. But $\alpha \rightarrow \alpha$ is a constraint in $Link(\mathcal{L})$ and $\alpha \rightarrow \beta$ a constraint in $Link(\mathcal{C}_2)$, and by *And*, $\langle a, a \rangle$ cannot be inappropriate in $LC_2 = Link(\mathcal{L} \parallel \mathcal{C}_2)$. But since $\tau \leq (\alpha \wedge \beta)$, $\langle a, a \rangle$ would indeed be externally inappropriate (wrt $\{a : \alpha\}$) in L_1, LC_2 and N if $\tau' \perp \gamma$ (using the complex definition of external inappropriateness).

This completes the proof of the theorem.

□

Bibliography

Aczel, P. (1988). *Non Well-Founded Set Theory.* CSLI Lecture Notes 14, Stanford: CSLI Publications.

Aczel, P. (1990). Replacement systems and the axiomatization of Situation Theory. In R. Cooper, K. Mukai, and J. Perry (Eds.), *Situation Theory and its Applications I.* CSLI Lecture Notes 22, Stanford: CSLI Publications.

Aczel, P. and R. Lunnon (1991). Universes and parameters. In J. Barwise, J. M. Gawron, G. Plotkin, and S. Tutiya (Eds.), *Situation Theory and its Applications II.* CSLI Lecture Notes 26, Stanford: CSLI Publications.

Adams, E. (1975). *The Logic of Conditionals.* Dordrecht: Reidel.

Agre, P. E. and D. Chapman (1987). Pengi: an implementation of a theory of activity. In *Proc. Sixth National Conference on Artificial Intelligence*, Seattle, WA.

Allen, J., J. Hendler, and A. Tate (Eds.) (1990). *Readings in Planning.* San Mateo, CA: Morgan Kaufmann.

Anderson, A. R. and N. D. Belnap (1975). *Entailment, Vol. 1.* Princeton, N.J.:Princeton University Press.

Asher, N. and M. Morreau (1991). Commonsense entailment: a modal theory of commonsense reasoning. In *Proc. Twelfth International Joint Conference on Artificial Intelligence*, Sydney, Australia.

Bacchus, F. (1989). A modest, but semantically well-founded, inheritance reasoner. In *Proc. Eleventh International Joint Conference on Artificial Intelligence*, Detroit, IL.

Barker, S. (1994). Unpublished manuscript. Dept. of Philosophy, University of Melbourne.

Barwise, J. (1986). Conditionals and conditional information. In J. Barwise (Ed.), *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1989a). Notes on branch points in Situation Theory. In J. Barwise (Ed.), *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1989b). *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1989c). Situations, facts, and true propositions. In J. Barwise (Ed.), *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1989d). Situations, sets and the axiom of foundation. In J. Barwise (Ed.), *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1989e). Unburdening the language of thought. In J. Barwise (Ed.), *The Situation in Logic.* CSLI Lecture Notes 17, Stanford: CSLI Publications.

Barwise, J. (1993). Constraints, channels and the flow of information. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters (Eds.), *Situation Theory and its Applications III.* CSLI Lecture Notes 37, Stanford: CSLI Publications.

Barwise, J. and R. Cooper (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy 4*(2).

Barwise, J. and R. Cooper (1991). Simple Situation Theory and its graphical representation. In J. Seligman (Ed.), *Partial and Dynamics Semantics III, DYANA Report R2.1.C.* Centre for Cognitive Science, University of Edinburgh.

Barwise, J. and J. Etchemendy (1990). Information, infons and inference. In R. Cooper, K. Mukai, and J. Perry (Eds.), *Situation Theory and its Applications I.* CSLI Lecture Notes 22, Stanford: CSLI Publications.

Barwise, J., D. Gabbay, and C. Hartonas (1994). Information flow and the Lambek calculus. Philosphy Department, Indiana Univerity. Available by ftp as `phil.indiana.edu:/pub/BarGabHar94.dvi`.

Barwise, J. and L. Moss (1991). Lectures on situation theory and its foundations. Unpublished manuscript, Philosphy Department, Indiana Univerity.

Barwise, J. and J. Perry (1983). *Situations and Attitudes.* Cambridge: MIT Press.

Barwise, J. and J. Seligman (1993). Imperfect information flow. In *Proc. Eighth IEEE Symposium on Logic in Computer Science.*

Barwise, J. and J. Seligman (1994). The rights and wrongs of natural regularity. In J. Tomberlin (Ed.), *Philosophical Perspectives, Vol. 8.* California: Ridgeview.

Bealer, G. (1982). *Quality and Concept.* Oxford: Clarendon.

van Benthem, J. (1994). A note on dynamic arrow logic. In J. van Eijck and A. Visser (Eds.), *Logic and Information Flow.* Cambridge: MIT Press.

Bouquet, P. and F. Giunchiglia (1994). Reasoning about theory adequacy: a new solution to the qualification problem. To appear in *Fundamenta Informaticae.* Also available by ftp as `ftp.mrg.dist.unige.it:pub/mrg-ftp/9406-13.ps.Z`.

Boutilier, C. (1989). A semantical approach to stable inheritance reasoning. In *Proc. Eleventh International Conference on Artificial Intelligence*, Detroit, IL.

Boutilier, C. (1992). *Conditional Logics for Default Reasoning and Belief Revision.* Ph. D. thesis, Department of Computer Science, Univerity of Toronto.

Brooks, R. A. (1986a). Achieving Artificial Intelligence through building robots. Technical Report AI Memo 899, MIT.

Brooks, R. A. (1986b). Intelligence without representation. Technical Report AI Memo, MIT.

Buvač, S., V. Buvač, and I. A. Mason (1995). Metamathematics of contexts. To appear in *Fundamenta Informaticae*. Also available via ftp as `ftp://sail.stanford.edu/pub/buvac/context.ps`.

Carlson, G. (1977). *Reference to Kinds in English.* Ph. D. thesis, University of Massachussetts, Amherst.

Cavedon, L. (1993). Conditional constraints and Default Logic. Unpublished manuscript. Centre for Cognitive Science, University of Edinburgh.

Cavedon, L. (1994). A Channel-Theoretic model for conditional logics (extended abstract). In *Conference on Information-Oriented Approaches to Logic, Language and Computation*, Moraga, CA. To appear in *CSLI Lecture Notes* series.

Cavedon, L. and S. Glasbey (1994). An information-flow model of generics (extended abstract). In *Fifth Symposium on Logic and Language*, Hungary. Full version to appear in *Acta Linguistica.*

Chellas, B. F. (1975). Basic conditional logic. *Journal of Philosophical Logic 4.*

Declerck, R. (1991). The origins of genericity. *Linguistics 29.*

Delgrande, J. P. (1988). An approach to default reasoning based on a first-order conditional logic. *Artificial Intelligence 36.*

Delgrande, J. P. (1992). Miscellaneous examples for default reasoning. Technical Report CMPT TR 92-05, Centre for Systems Science, Simon Fraser University.

Devlin, K. (1991). *Logic and Information.* Cambridge University Press.

Dretske, F. (1981). *Knowledge and the Flow of Information.* Cambridge: MIT Press.

Dretske, F. (1988). *Explaining Behaviour.* Cambridge: MIT Press.

Dunn, J. M. (1984). Relevance logic and entailment. In D. Gabbay and F. Guenthner (Eds.), *Handbook of Philosophical Logic, Volume III.* Dordrecht: Reidel.

Escriba, J. B. (1992). Two formal systems for Situation Semantics. *Notre Dame Journal of Formal Logic 33*(1).

Etherington, D. W. (1988). *Reasoning with Incomplete Information.* London: Pitman.

Etherington, D. W., S. Kraus, and D. Perlis (1991). Nonmonotonicity and the scope of reasoning. *Artificial Intelligence 52.*

Fenstad, J. E., P.-K. Halvorsen, T. Langholm, and J. van Benthem (1987). *Situations, Language and Logic.* Dordrecht: Reidel.

Fodor, J. (1990). *A Theory of Content and Other Essays.* Cambridge: MIT Press.

Freund, M. and D. Lehmann (1993). Nonmonotonic inference operations. *Bulletin of the Interest Group in Pure and Applied Logic 1*(1).

Freund, M., D. Lehmann, and P. Morris (1991). Rationality, transitivity and contraposition. *Artificial Intelligence 52.*

Gabbay, D. M. (1985). Theoretical foundations for non-monotonic reasoning in expert systems. In K. R. Apt (Ed.), *Proc. of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems.* Springer-Verlag.

Gardenfors, P. (1988). *Knowledge in Flux.* Cambridge: MIT Press.

Gardenfors, P. (1993). Conceptual spaces as a basis for cognitive semantics. Technical report, Philosophy Department, Lund University.

Georgeff, M. P. (1987a). Actions, processes and causality. In M. P. Georgeff and A. Lamsky (Eds.), *Reasoning about Actions and Plans: Proceedings of the 1986 Workshop*. Los Altos, CA: Morgan Kaufman.

Georgeff, M. P. (1987b). Planning. In *Annual Review of Computer Science*. Annual Reviews Inc.

Georgeff, M. P., D. N. Morley, and A. Rao (1993). Events and processes in Situation Semantics. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters (Eds.), *Situation Theory and its Applications III*. CSLI Lecture Notes 37, Stanford: CSLI Publications.

Ginsberg, M. L. and D. E. Smith (1988). Reasoning about action II: the qualification problem. *Artificial Intelligence 35*.

Girard, J.-Y. (1987). Linear logic. *Theoretical Computer Science 50*.

Giunchiglia, E. and P. Traverso (1993). A multi-context architecture for formalizing complex reasoning. To appear in *International Journal of Intelligent Systems*. Also available as Tech. Report 9307-26, IRST, Trento, Italy.

Giunchiglia, F., L. Serafino, E. Giunchiglia, and M. Frixone (1993). Non-omniscient belief as context-based reasoning. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence*, Chambery, France.

Groenendijk, J. and M. Stokhof (1991). Dynamic predicate logic. *Linguistics and Philosophy 14*(1).

Guha, R. V. (1991). *Contexts: A Formalization and Some Applications.* Ph. D. thesis, Stanford University. Also available as MCC Technical Report ACT-CYC-423-91.

Halpern, J. and Y. Moses (1985). Guide to modal logics of knowledge and belief: preliminary report. In *Proc. Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA.

Healey, P. and C. Vogel (1994). A Situation Theoretic model of dialogue. In K. Jokinen (Ed.), *Pragmatics in Dialogue Management.* Gothenburg Monographs in Linguistics.

Healey, P., C. Vogel, M. Paul, and L. Cavedon (1993). Ontological pluralism in dialogue (extended abstract). In *Proceedings of the Third International Colloquium on Cognitive Science*, San Sebastian, Spain.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases.* Ph. D. thesis, Dept. of Linguistics, University of Massachusetts at Amherst.

Hintikka, J. (1962). *Knowledge and Belief.* Ithaca, N.Y.: Cornell University Press.

Horty, J., R. Thomason, and D. Touretzky (1990). A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence 42*.

Horty, J. F. and R. H. Thomason (1988). Mixing strict and defeasible inheritance. In *Proc. Seventh National Conference on Artificial Intelligence.*

Hunter, G. (1980). Conditionals, indicative and subjunctive. In J. Dancey (Ed.), *Papers on Language and Logic.* Keele University Library.

Israel, D. and J. Perry (1990). What is information. In P. P. Hanson (Ed.), *Information, Language and Cognition*. Vancouver, B.C: University of British Columbia Press. Also available as report no. CSLI-91-145, CSLI Publications, Stanford.

Israel, D., J. Perry, and S. Tutiya (1993). Executions, motivations and accomplishments. *Philosophical Review 4*.

Jackson, F. (1987). *Conditionals*. Oxford: Basil Blackwell.

Kamp, H. (1981). A theory of truth and semantic representation. In J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof (Eds.), *Formal Methods in the Study of Language*. Amsterdam: Mathematische Centrum.

Kamp, H. and U. Reyle (1993). *From Discourse to Logic*. Studies in Linguistics and Philosophy. Dordrecht: Reidel.

Konolige, K. (1988). On the relationship between default and autoepistemic logic. *Artificial Intelligence 35*.

Konolige, K. (1989). Hierarchic autoepistemic theories for nonmonotonic reasoning: preliminary report. In *Non-Monotonic Reasoning*. Springer-Verlag. Lecture Notes in Artificial Intelligence 346.

Koons, R. C. (1993). Nonmonotonic projection, causation and induction. In P. Aczel, D. Israel, Y. Katagiri, and S. Peters (Eds.), *Situation Theory and its Applications III*. CSLI Lecture Notes 37, Stanford: CSLI Publications.

Koons, R. C. (1994). Information, representation and the possibility of error. In *Conference on Information-Oriented Approaches to Language, Logic and Computation*, Moraga, CA. To appear in *CSLI Lecture Notes* series.

Kraus, S., D. Lehmann, and M. Magidor (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence 44*(1).

Krifka, M. (1987). An outline of genericity. Technical Report SNS-Bericht 87-25, Seminar für Natürlich-sprachliche Systeme, Tübingen University. (Partly in collaboration with Claudia Gerstner.).

Krifka, M. (1995). Focus and the interpretation of generic sentences. In F. Pelletier and G. Carlson (Eds.), *The Generic(s) Book*. Chicago University Press. To appear.

Landman, F. (1986). *Towards a Theory of Information*. Dordrecht: Foris.

Lehmann, D. (1989). What does a conditional knowledge base entail? In *Proc. First International Conference on Knowledge Representation and Reasoning*, Toronto, Canada.

Lemon, O. J. (1993). Perspectives on relevance: The logic of mental spaces. Master's thesis, Centre for Cognitive Science, University of Edinburgh.

Lesperance, Y. and H. J. Levesque (1995). Indexical knowledge and robot action: a logical account. *Artificial Intelligence 72*(1–2). Special edition on "Computational Theories of Interaction and Agency", Agre, P. E. and Rosenschein, S. J. (eds).

Levesque, H. J. (1984). A logic of implicit and explicit belief. In *Proc. National Conference on Artificial Intelligence*, Austin, TX.

Lewis, D. (1973). *Counterfactuals*. Cambridge, MA: Harvard.

Lewis, D. (1986). Causation. In D. Lewis (Ed.), *Philosophical Papers, Vol. II.* Oxford University Press.

Lifschitz, V. (1985). Computing circumscription. In *Proc. Ninth International Conference on Artificial Intelligence*, Los Angeles, CA.

Lifschitz, V. (1987). Formal theories of action. In F. Brown (Ed.), *The Frame Problem in Artificial Intelligence.* Los Altos, CA: Morgan Kaufman.

Lifschitz, V. (1989). Benchmark problems for formal nonmonotonic theories. In Reinfrank, de Kleer, Ginsburg, and Sandewall (Eds.), *Workshop on Non-Monotonic Reasoning.* Springer-Verlag. Lecture Notes in Artificial Intelligence 346.

Lifschitz, V. (1990). Frames in the space of situations. *Artificial Intelligence 46.*

Lukasziwicz, W. (1990). *Non-Monotonic Logic: Formalisation of Commonsense Reasoning.* Ellis Horwood.

Makinson, D. (1989). General theory of cumulative inference. In Reinfrank, de Kleer, Ginsburg, and Sandewall (Eds.), *Workshop on Non-Monotonic Reasoning.* Springer-Verlag. Lecture Notes in Artificial Intelligence 346.

Mares, E. D. and A. Fuhrmann (1993). A relevant theory of conditionals. To appear in *Journal of Philosophical Logic.*

McCarthy, J. (1980). Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence 13.*

McCarthy, J. (1986). Applications of circumscription to formalizing commonsense reasoning. *Artificial Intelligence 28.*

McCarthy, J. (1993). Notes on formalizing context. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence*, Chambery, France.

McCarthy, J. M. and P. J. Hayes (1969). Some philosophical problems from the standpoint of Artificial Intelligence. In Meltzer and Michie (Eds.), *Machine Intelligence.* Edinburgh University Press.

ter Meulen, A. (1986). Generic information, conditional contexts and constraints. In E. C. Traugott, A. ter Meulen, J. S. Reilly, and C. A. Ferguson (Eds.), *On Conditionals.* Cambridge University Press.

Moore, R. C. (1984). Possible-world semantics for autoepistemic logic. In *Proceedings of Non-Monotonic Reasoning Workshop*, New Paltz, NY.

Moore, R. C. (1985). Semantical considerations on nonmonotonic logic. *Artificial Intelligence 25.*

Morreau, M. (1992a). *Conditionals in Philosophy and Artificial Intelligence.* Ph. D. thesis, Department of Philosophy, University of Amsterdam.

Morreau, M. (1992b). Norms or inference tickets? In *Working Papers of the Fourth International Workshop on Nonmonotonic Reasoning.* Also available in Morreau and Veltman (eds.), *Commonsense Entailment and Update Semantics,* DYANA Report R2.5.C, Centre for Cognitive Science, University of Edinburgh.

Muskens, R. A. (1989). Going partial in montague grammar. In R. Bartsch, J. F. A. K. van Benthem, and P. van Emde Boas (Eds.), *Semantics and Contextual Expressions: Proceedings of the Sixth Amsterdam Colloquium.* Dordrecht: Foris.

Nivre, J. (1992). *Situations, Meaning and Communication.* Ph. D. thesis, Department of Linguistics, University of Gothenburg.

Nute, D. (1980). *Topics in Conditional Logic.* Dordrecht: Reidel.

Nute, D. (1984). Conditional logic. In D. Gabbay and F. Guenthner (Eds.), *Handbook of Philosophical Logic, Vol. II.* Dordrecht: Reidel.

Nute, D. (1993). Defeasible Prolog. Technical Report AI-1993-04, AI Programs, University of Georgia. Presented at AAAI Fall Symposium on Automated Deduction in Nonstandard Logics, Raleigh, NC, 1993.

Padgham, L. (1989). Negative reasoning using inheritance. In *Proc. Eleventh International Joint Conference on Artificial Intelligence*, Detroit, IL.

Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* San Mateo: Morgan Kaufman.

Pearl, J. (1990). System Z: A natural ordering of defaults with tractable applications to default reasoning. In M. Vardi (Ed.), *Proc. Theoretical Aspects of Reasoning about Knowledge.* San Mateo: Morgan Kaufman.

Pelletier, F. and G. Carlson (Eds.) (1995). *The Generic(s) Book.* Chicago University Press. To appear.

Pollock, J. (1976). *Subjunctive Reasoning.* Dordrecht: Reidel.

Pollock, J. L. (1987). Defeasible reasoning. *Cognitive Science 11.*

Read, S. (1988). *Relevant Logic: a Philosophical Examination of Inference.* Oxford: Basil Blackwell.

Reiter, R. (1980). A logic for default theory. *Artificial Intelligence 13.*

Restall, G. A. (1994). Information flow and relevant logics. In *Conference on Information-Oriented Approaches to Logic, Language and Computation*, Moraga, CA. To appear in *CSLI Lecture Notes* series.

Rosenschein, S. J. (1985). Formal theories of knowledge in AI and robotics. *New Generation Computing 3*(4).

Rosenschein, S. J. and L. P. Kaelbling (1986). The synthesis of machines with provable epistemic properties. In J. Halpern (Ed.), *Theoretical Aspects of Reasoning about Knowledge: Proceedings of the 1986 Conference.* Los Altos: Morgan Kaufman.

Rounds, W. C. and G.-Q. Zhang (1995). Domain theory meets default logic. *Journal of Logic and Computation 5*(1).

Routley, F. R. and R. K. Meyer (1973). The semantics of entailment (I). In H. Leblanc (Ed.), *Truth, Syntax and Modality.* Amsterdam: North Holland.

Ryan, M. (1992). *Ordered Presentations of Theories—Default Reasoning and Belief Revision.* Ph. D. thesis, Department of Computing, Imperial College.

Seligman, J. (1990). Perspectives in Situation Theory. In R. Cooper, K. Mukai, and J. Perry (Eds.), *Situation Theory and its Applications I.* CSLI Lecture Notes 22, Stanford: CSLI Publications.

Seligman, J. (1993). The Structure of Regularity: unpublished lecture notes on Channel Theory. Available by ftp as `phil.indiana.edu:/pub/SORnotes/*`.

Seligman, J. and J. Barwise (1993). Channel Theory: Towards a Mathematics of Imperfect Information Flow (draft). Draft available by ftp as `phil.indiana.edu:/pub/SelBar93.ps`.

Shannon, C. E. and W. Weaver (1949). *The Mathematical Theory of Communication.* Urbana: University of Illinois Press.

Shoham, Y. (1987). A semantical approach to nonmonotonic logics. In *Proc. IEEE Symposium on Logic in Computer Science*, Ithaca, NY.

Shoham, Y. (1988). *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence.* Cambridge: MIT Press.

Smith, B. C. (1987). Semantics of clocks. Technical Report CSLI-87-100, CSLI.

Smith, B. C. (1992). The owl and the electronic encyclopedia. In D. Kirsh (Ed.), *Foundations of Artificial Intelligence.* Cambridge: MIT Press.

Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in Logical Theory.* Oxford: Basil Blackwell.

Stalnaker, R. (1987). *Inquiry.* Cambridge: MIT Press.

Stein, L. A. (1992). Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence 55.*

Suchman, L. A. (1987). *Plans and Situated Actions.* Cambridge University Press.

Touretzky, D. (1986). *The Mathematics of Inheritance Systems.* London: Pitman.

Touretzky, D., J. Horty, and R. Thomason (1987). A clash of intuitions: the current state of nonmonotonic inheritance systems. In *Proc. Tenth International Joint Conference on Artificial Intelligence*, Milan, Italy.

Touretzky, D. S. (1984). Implicit ordering of defaults in inheritance systems. In *Proc. National Conference on Artificial Intelligence*, Austin, TX.

Turner, R. (1987). A theory of properties. *Journal of Symbolic Logic 52*(2).

Veltman, F. (1993). Defaults in update semantics. To appear in *Journal of Philosophical Logic.*

Vogel, C. (1995). *Inheritance Reasoning: Psychological Plausibility, Proof Theory and Semantics.* Ph. D. thesis, Centre for Cognitive Science, University of Edinburgh.

Winograd, T. and C. F. Flores (1986). *Understanding Computers and Cognition.* NJ: Ablex Corp.

Wobcke, W. (1995). A theory of conditionals based on hierarchies of situations. To appear in *Cognitive Science.*

Zhang, G.-Q. and W. C. Rounds (1993). Nonmonotonic consequences of default model theory. In *Proc. Bar-Ilan Symposium on Mathematics and Artificial Intelligence*, Israel.