

"Making Moral Space: A Reply To Churchland" CANADIAN JOURNAL OF PHILOSOPHY
(MORAL EPISTEMOLOGY NATURALIZED
Supplementary Volume for the Canadian Journal of Philosophy
Edited by Richmond Campbell and Bruce Hunter)

Making Moral Space: A Reply To Churchland

Andy Clark
School of Cognitive and Computing Sciences
University of Sussex
Brighton, BN1 9QH
UK

Like those famous nations divided by a single tongue, my paper (this volume) and Professor P.M. Churchland's deep and engaging reply offer different spins on a common heritage. The common heritage is, of course, a connectionist vision of the inner neural economy- a vision which depicts that economy in terms of supra-sentential state spaces, vector-to-vector transformations, and the kinds of skillful pattern-recognition routine we share with the bulk of terrestrial intelligent life-forms. That which divides us is, as ever, much harder to isolate and name. Clearly, it has something to do with the role of moral talk and exchange, and something to do with the conception of morality itself (and, correlatively, with the conception of moral progress). Most of this Reply will be devoted to clarifying the nature of the disputed territory. First, though (as a prophylactic against misunderstanding) I shall rehearse some points of agreement concerning moral talk and progress.

Professor Churchland and I agree that words, talk, moral labels, and the whole collective infrastructure of moral texts, rules, traditions, tools and practices matter. We agree, indeed, that it is this species-specific overlay that gives human thought and reason (in science, morals, and elsewhere) much of its distinctive power and character. Thus Churchland recognizes and emphasizes "the genuine novelty represented by the evolutionary emergence of language and the cultural emergence of discursive rules", and depicts this novelty as something that "extends the reach and elevates the quality of the original non discursive cognition". Furthermore, we agree that words, texts and technologies are, taken alone, cognitively and morally inert, so that "even when that external machinery does get deployed, it is the original and more basic form of cognition that does the deploying. Rules are useless unless the capacity for reliable perception of their categories is already in place, and such perception depends utterly on the inarticulable processes of vector coding and prototype activation". We agree also that even our best moral

rules, maxims and guidelines (even God's own moral rules, maxims and guidelines, were God forced to formulate them as sentences in some Earthly tongue) are not to be viewed as displaying the full, rich content of our own (or God's) achieved moral expertise. Instead, the rules, maxims and guidelines play a kind of facilitating role. They act as reference points for collaborative moral reasoning and discussion, and they sow the seeds for deeper and more penetrating moral thought (for example, by providing summary labels which can support the discovery of deeper, more abstract, otherwise cognitively invisible, moral concepts- as per the discussion of pan troglodytes in my text). And we agree, finally (though here I suspect I was not clear in the text) that there can be, and indeed has been, genuine moral progress. I concur wholeheartedly with professor Churchland's forceful description of the many moral mistakes and attitudes enshrined in the Bible and believe, as he does, that many contemporary moral norms represent substantial advances over that primitive foray into moral space.

Where, then, do we disagree? We disagree, I suggest, on two (related) counts.

First, we disagree on the precise role of all that external scaffolding and moral infrastructure. As Professor Churchland _as it, the role of the scaffolding is largely to offload, preserve, stockpile and share our collective moral wisdom and experience. And moral wisdom itself is conceived as a kind of know-how concerning the successful navigation of social space, a type of know-how we thus share with many other social animals including "baboon troops, wolf packs, dolphin schools, chimpanzee groups, lion prides". What we find in such cases, Churchland suggests, is "the same complex ebb and flow of thoughtful sharing, mutual defense, fair competition, familial sacrifice, staunch alliance, minor deception, major treachery, and the occasional outright ostracism that we see displayed in human societies."

What we do not find, he notes, is the peculiar kind of discursive language-use or highly articulated non-biological infrastructure that characterizes human societies. As a result, in the case of other social animals "their social cognition is conducted entirely within the more primitive and nondiscursive form of cognition". The specific social spaces we might navigate are, Churchland allows, deeply transformed by these extra layers of infrastructure. But the discursive infrastructures, Churchland insists, "do not bring moral reasoning into existence for the first time, and they do not provide a conceptual model remotely adequate to the phenomenon of moral cognition in single individuals and nonhuman animals" Moral understanding, it seems, is a more primitive thing.

Here, then, is the first point of real disagreement. For on my account, our practices of moral talk and exchange, and our collective efforts to create the kinds of abstract, shared conceptions (of 'charity' 'rights' 'equality' 'opportunity' etc.) that such discussions require are part of what constitutes our practices as genuinely moral in the first place. I do not dispute, in any way, Professor Churchland's depiction of nonlinguistic animals as navigating social spaces. But I do dispute the apparent direct assimilation of such skilled navigation to moral activity. There is, I maintain, a sufficiently profound difference between our human moral projects and the project of successful social navigation to justify treating the latter, but not the former, as distinctively moral modes of thought and reason. Such modes are marked, for example, by the requirement to provide reasons for our actions, and to be able to address the important question of the acceptability, or otherwise, of our own underlying needs, desires and goals. They are marked also, I argued, by an essential commitment to collaborative moral endeavor: to finding routes through moral space that accommodate multiple perspectives and points of view. Practices of public moral discussion and exchange creates, I tried to argue, these kinds of moral sensitivity in much the same way as the creation of financial institutions creates the space to trade in stocks, shares, options and futures, then options on futures, and so on. In each case the presence of the infrastructure is partly constitutive of the very possibility of the target phenomenon.

In weak support of this rather strong thesis (the thesis, if you like, of the discursive construction of moral space) I offered a couple of more detailed- but admittedly non-moral- exemplars. One involved the ability of label-exploiting chimps (pan troglodytes) to grasp kinds of abstraction beyond the ken of their un-augmented cousins. The other involved the role of numerals in enabling our communal exploration of mathematical space. The idea here, which I should have made more explicit, was that despite the clear overlap in base-line neural skills, only the activity of the numeral-enhanced humans counts as genuinely mathematical. The very concept of a number, I would argue, is available to our species only courtesy of its experiences with the artifactual domain of numeral construction and manipulation. What other animals do by way of counting is not properly conceived as mathematics in the absence of that overlay, just as their skills at social navigation, in the absence of a similar overlay of discursive concepts, do not count as moral skills.

The second point of disagreement follows rather directly from the first. While we both agree on the possibility of moral progress, we harbor subtly different visions of in what such progress might consist. For Professor Churchland progress consists in increased collective success at the negotiation of increasingly complex social spaces. I claim, by contrast (at least I think it is by contrast-see below) that moral progress consists primarily in increased collective sensitivity to the needs, reasons and desires of others. Our communal explorations of moral space serve to sculpt and tweak these needs and desires while simultaneously attempting to accommodate as wide a variety as possible. Now in practice, I concede (hence my hesitation above) that it will be hard indeed to distinguish Churchland's vision from mine. For the typical upshot of all this mutual consideration of needs, reasons and arguments should, one hopes, be a smoother, gentler social swirl. But the emphasis (on the exchange of reasons versus the navigation of social space) strikes me as important. For we make moral progress, I want to claim, only by swimming better in a sea of other's needs and reasons, not by simply swimming better in a social sea.

This difference in the conception of the moral domain explains, I think, some of my continued resistance to professor Churchland's radical visionary stance concerning the future of folk-psychology. While agreeing that future human brains may well come to deploy new and better modes of thought and reason (for example, by learning, courtesy of games such as SIMCITY, better ways to think about complex, decentralized, self-organizing phenomena) I find myself unable to conceive of the future morality that by-passes the communal exchange of discursive representations. And I lack a conception of in what a post-sentential exchange of reasons and justifications might consist (even using a diagram to make a point often depends on some accompanying sentential gloss). I am happy to concede, however, that my failures of imagination are just that, and no more (they are not (precisely!) arguments). So I now record an open verdict on the possible forms of future moral exchange and debate, while still insisting that there must be such exchange and debate on pain of failing to re-constitute any genuinely moral realm. Here to stay, I claim, must be some form of interpersonal discursive representation capable of providing rough summary abstractions of the rich contents encoded in high dimensional state spaces. Such abstractions, I argue, play vital roles both in learning and in collaborative thought. In the case at hand, such representations do not simply oil the wheels of moral debate, they actively constitute the thinking as moral.

More generally, Philosophers of Cognitive Science (with the notable exception of Dan Dennett - (see e.g. Dennett (1996)), tend to underestimate just how very special we humans are. This downplaying is doubtless the result of an otherwise laudable desire to keep things natural and to emphasize the deep and real continuities between human cognition and that of other animals. But we are different, and the difference is cognitively deep (even if rooted in only some small neural difference). To appreciate the difference we must abandon our staunchly brain-and-individual oriented stance, and attend equally to the potent cognitive transformations effected by the matrix of words and technologies in which we live and think.

Common ground thus marked, and disputed territory highlighted, what is to be done? Here, I confess, I am at something of a loss. For all I have done, on reflection, is to present a personal, biased picture of in what moral cognition might consist, and to accompany this picture with a couple of (notably non-moral) illustrations. The picture is one in which the moral realm comes into view, and moral cognition is partially constituted, only by the joint action of neural resources we share with other animals and the distinctively human infrastructure of linguaform moral debate and reason.

Our status as moral agents depends crucially, if I am right, on the many additional layers of cognitive circuitry we have slowly woven into the worlds within which we now think, reason, act, build and legislate. But Professor Churchland's vision, so wonderfully expressed and powerfully argued in his contribution, stands out as clear and compelling. The moral realm, as he depicts it, is one already explored by many social animals, and is not at all the peculiar province of the language-and-culture enhanced ('mindware upgraded'- see Clark, In Press) human species. Who (if either) is right? And how can we tell? My closing thought is that this is, in all likelihood, not exactly an empirical question. The answer depends upon some hard decisions concerning which aspects of current moral practice should be foregrounded in our best philosophical and scientific treatments of morality. And that, I venture to suggest, may be a moral, rather than a properly scientific, question.

References

Churchland, P.M. (This Volume) "Rules, Know-How, And The Future of Moral Cognition"

Clark, A (This Volume) "Word and Action: Reconciling Rules and Know-How In Moral Cognition"

Clark, A (In Press) Mindware: An Introduction To The Philosophy Of Cognitive Science (Oxford University press).