

To appear in *Oxford Handbook of Philosophy and Neuroscience*,
edited by John Bickle, (2008) Oxford University Press.

Monday, January 22, 2007

INFERENCE TO THE BEST DECISION¹

Patricia Smith Churchland

Philosophy Department UCSD

La Jolla CA 92093

Anyone who is tempted to explore the possibility of neuroethics – the idea that what is and is not valuable is rooted in basic biology – can expect to be scolded with dictum:

“you cannot derive an *ought* from an *is*”.² For those of us who recognize evolutionary

¹ The basic idea for this article can be found in *Natural Ethical Facts* (Cambridge: MIT, 2001), by William Casebeer, wherein he observed that many judgments about what ought to be done are abductive, and in Paul Churchland’s “Toward a cognitive neurobiology of the moral virtues” *Topoi* 17 (1998): 83-96 wherein he shows how parameter-space representation yields prototypes and a similarity metric. Parameter-space representation is the epistemological bridge to inference to the best decision. See also Lakoff, *Moral Politics*. (Chicago:University of Chicago Press,1996) and Mark L. Johnson “Ethics”, in: *A Companion To Cognitive Science*, ed. by W. Bechtel and G. Graham (Oxford: Blackwells, 1998): 691-701. I also got sensible advice from Dan Dennett, David Brink, John Jacobson, Dale Dorsey, and Michael Stack.

² For a recent and uncompromising defense of Hume’s claim, see Philip Kitcher, “Biology and Ethics”, *The Oxford Handbook of Ethics* (Oxford: Oxford University Press, 2006) Ed. by David Copp. Oxford (forthcoming). See also Catherine Wilson “The biological basis and ideational superstructure of ethics.” In: *Moral Epistemology*

biology as the essential backdrop for inquiries into human nature and human behavior, the assumption that this dictum has axiomatic status looks increasingly problematic. Morality is not the product of a mythical pure reason divorced from natural selection and the neural wiring that motivates the animal to sociability. It emerges from the human brain and its responses to real human needs, desires, and social experience; it depends on innate emotional responses, on reward circuitry that allows pleasure and fear to be associated with certain conditions, on cortical networks, hormones and neuropeptides. Its cognitive underpinnings owe more to case-based reasoning than to conformity to rules.

That we are social animals with social dispositions is a central fact owed to our evolutionary history.³ Sociability has been selected for in humans, as well as in baboons, wolves, ravens, jays, dolphins, chimpanzees and many other species. Oxytocin, known to play a role in parturition and parent-offspring interactions, also plays a role in affiliative behaviors such as sex and grooming.⁴ Vasopressin also plays a crucial role, especially in males. As always in biology, there are individual differences within a species, and certainly one can observe individual differences among humans with respect to social dispositions and the capacity to learn what is expected in our own social group; for

Naturalized, edited by R. Campbell and B. Hunter, *Canadian Journal of Philosophy Supplementary* Vol. 26, (2000): 211-244.

³ John M. Allman, *Evolving Brains* (New York: Scientific American Library, 1999); R.I.M. Dunbar, "The social brain hypothesis". *Evolutionary Anthropology* Vol. 6 (1998): 178-90.

⁴ For a comprehensive review article, see Thomas R. Insel and Russel D. Fernald "How the brain processes social information: Searching for the social brain." *Annual Review of Neuroscience* Vol. 27 (2004): 697-722. See also Frances P. Champagne and James P. Curley How social experiences influence the brain. *Current Opinion in Neurobiology* 15 (2005):704-709.

example, in the case of MAOA variants.⁵ There are also individual differences in temperamental features such as risk aversion, sensitivity to disapproval and capacities for self-control.

At a neurobiological level, we are beginning to understand how sociability is supported in the brain, and the role of experience in learning group standards of behavior. We are also beginning to understand the relationship between gene expression and epigenetic factors, and their impact on socially appropriate behavior in maturity.⁶ For example, research shows that infant cuddling (licking, in rats) initiates a cascade of neurochemical events that eventually alters gene expression that modifies the circuitry mediating social capacities. At a more general level, biologically constrained models demonstrate how traits of cooperation and social orderliness can spread through a population; how the virtues can be a benefit, cheating a cost, and punishment of the socially dangerous a necessity.⁷ All this is consistent with natural selection and in no way implies group selection. In short, owing to developments over the past three decades, a tension has developed between the sanctity of the “ought/is” dogma and what is known about the neurobiology of social behavior. As I shall argue, the cognitive process that we loosely

⁵ The MAOA (monoaminoxidase-A) gene is carried on the X chromosome and some individuals carry a mutation that results in decreased expression of the gene. MAOA is an enzyme that metabolizes serotonin and norepinephrine. If the individual carrying the MAOA is male, then the probability of his displaying irrational and self-destructive violent behavior is high, and very high if, as a child, he is abused. See A. Caspi et al. “Role of genotype in the cycle of violence in maltreated children.” *Science* 297 (2002): 851-854.

⁶ I. Weaver, N. Cervoni, F. Champagne, A. D’Allesio, S. Sharma, J. Seckl, M. Szyf and M. Meaney. Epigenetic programming by maternal behavior. *Nature Neuroscience*. 7 (2004): 847-854.

⁷ For an accessible overview of the literature, see Matt Ridley, *The Origins of Virtue: Human Instinct and the Evolution of Cooperation*. 1996. Viking.

call *inference to the best decision* is a solution to this tension. But first, I shall consider the epistemological background.

The idea of “inference to the best *explanation*”, often referred to by philosophers as “abduction”, is also known in experimental psychology as “case-based reasoning”.⁸ Essentially, case-based reasoning yields a solution to a problem (what is this, how does this work, why did this happen) by using memory for relevantly similar cases, and applying past knowledge to present circumstances. The point I shall emphasize below, is that case-based reasoning, whether used for addressing scientific or moral questions, does not rely on universal propositions – neither laws, in the domain of science, nor maxims nor moral theories nor moral rules, in the domain of social behavior.⁹ What it does rely on are prototypes, similarity metrics, and analogies.¹⁰ Moreover, neural network models demonstrate that networks easily learn from examples and the response patterns of the inner units display a similarity metric. The parameter spaces the inner units represent during training, even unsupervised training, are in fact similarity spaces.

In philosophy, abduction has been embraced as a solution to a problem in epistemology. The problem was this: according to one plausible theory -- the deductive-nomological (D-N) theory of explanation -- acceptable explanations of a phenomenon are deductive arguments which must contain nomological statements (natural laws) of the form, “All

⁸ David B. Leake. “Case-based reasoning”. In: *A Companion to Cognitive Science*, ed by W. Bechtel and G. Graham. (Oxford: Blackwells 1998): 465-476. There does exist an extensive research program into the nature of case-based reasoning.

⁹ See Paul Churchland op.cit

¹⁰ See George Lakoff ibid and Mark L. Johnson ibid.

X's are Y's". An illustration of where the theory works is the following: why did that copper expand? Because (Premise 1) copper expands when heated (nomological statement) and (Premise 2) that piece of copper was heated. Therefore, (3) that piece of copper expanded.

The trouble was, in actual scientific practice, as well as in the day-to-day business of life, people routinely generate good, powerful and highly predictive explanations without relying on natural laws or other generalizations in their explanations. The D-N theory of explanation, while apparently plausible, turned out to be deeply problematic for many exemplary scientific explanations. For example, consider the explanation of how information is coded in DNA; the explanation of the origin of the earth, the origin of species, the function of the heart and lungs, how the pancreas works, why the dinosaurs became extinct, the cause of northern lights, the tides, and cervical cancer. Moreover, many routine, commonplace explanations likewise fail to conform to the requirement that a universal statement be included in the premises. They are good explanations in the sense that they launch powerful predictions that turn out to be correct and they inspire manipulations that turn out to be successful. But they do not depend on universal propositions, and they do not involve derivation in any straightforward logical sense.

How do these nomically-impooverished explanations nonetheless succeed in explaining? As Peirce argued, and as we all know now, empirical understanding is mainly acquired by recognizing an event as relevantly similar to a familiar class and inferring that the event has a cluster of properties similar to that of the class. Abduction relies on the capacities to generalize usefully from observed cases and draw suitable analogies from

the familiar to the unfamiliar.¹¹ Thus abductive – case-based --explanations address mechanism, origins, causal organization, and so forth.¹² Much of the cognitive business of abduction, like the cognitive business of perception and behavioral control, is probably nonconscious and largely inarticulable. In some instances, case-based reasoning may reach for rather abstract analogies between the familiar and the puzzling; other instances may be fairly humdrum. As Peirce noted, abduction sometimes amounts to sophisticated perceptual pattern recognition, as when the neurologist instantly recognizes a tremor as a sign of Parkinson’s disease, or an astronomer recognizes a fuzzy spiral in the night sky as a distant galaxy.¹³ Sometimes they involve quite abstract analogies, as when Newton realized that the revolution of the moon was, at bottom, like falling towards a large gravitational mass,¹⁴ and when Darwin realized that the origin of species was essentially like breeding dogs but without the breeder.

Why does it hurt, screams the child, and the mother, given her background understanding (it is summer and wasps are about) and her observation of the growing red welt, along with her own similar experience in similar conditions, infers that the child has been stung

¹¹ David Leake, editor. *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. (Cambridge:AAAI Press/MIT Press: 2002). Steve Bogaerts and David Leake (2005) Facilitating CBR for incompletely described cases: distance metrics for partial problem descriptions.

¹² For a discussion of how children acquire causal understanding, see C. Glymour, A. Gopnik, D.M.Sobel, L. E. Schultz, T. Kushnir and D. Danks. “A theory of causal learning in children: Causal maps and Bayes nets.” *Psychological Review*, 111 (2004): 1-31. See Carl Craver, “Beyond Reduction: mechanisms, Multifield Integration, and the Unity of Science”, *Studies in History and Philosophy of Biological and Biomedical Sciences*.(2005)

¹³ See Paul Churchland *ibid*.

¹⁴ Paul Churchland *The Engine of Reason, The Seat of the Soul*, (Cambridge:MIT Press: 1995).

by a wasp. Inference to the best explanation. She need not invoke, implicitly or otherwise, a nomic statement about red welts in general. To simplify, she recognizes this red welt as relevantly like others she has seen -- welts that were caused by wasp stings. In short, *pattern recognition* rather than *discursive argument* is the essential cognitive platform.

The space shuttle Challenger crashes, the scientific committee investigates. Many possible explanations are ruled out. Feynman puts sample O-rings (gaskets) in ice-water, observes the onset of brittleness. Probably the cold temperature of the o-rings was the critical factor that precipitated the catastrophe. Inference to the best explanation of the explosion. Watson and Crick, using background knowledge in chemistry and physics, and mulling over the X-Ray photograph of the DNA crystal, build a model of how the DNA molecule might be structured. They realize that the four base-pair organization of the inner ladder is like a code, and the double helix can split and reform to replicate itself. Inference to the best (albeit, incomplete) explanation of a mechanism whereby copying of information from parent to offspring is achieved. Empirical science being what it is, the inferred explanations are only probable, and may be revised in the light of new understanding, but case-based explanation is often the best one can do, relative to the conceptual and evidential resources available in the time available. In addition, because scientific data are always partial and understanding may vary when conceptual frameworks vary, there may well be disagreements that cannot be resolved, at least in the short run. It is also significant that no one argues for the need for a background rule for

prudential *oughts*. So if we can infer what prudentially we ought to do, without aid of a prudential background rule, why not in the moral domain?

Drawing on evolutionary biology, experimental psychology and neurobiology, we can approach case-based reasoning as a brain phenomenon. Nervous systems were selected for because they allow an organism to move, rather than passively take what comes.¹⁵ The fundamental functions of nervous systems are to move, survive, and to make predictions that inform movement, thereby enhancing the organism's chances of surviving long enough to reproduce. Other things being equal, an organism that can predict events – where good food and shelter are, whether it is best to run or hide, whether this is a good mate – will gain in the competition to survive and reproduce. Behaviorally useful categorizations of relevantly similar events, and the retrieval capacity swiftly to access a category for use when needed, are fundamental to prediction and hence to survival.

Case-based reasoning, whatever exactly the neural mechanisms involved, is rooted in neurobiological dispositions to categorize “me-relevant” stimuli for the purposes of prediction that will guide behavior, either in the short or the long run. Some of this understanding may be nonconscious, much of it is undoubtedly organized in prototype mode rather than in propositional mode, especially for animals with no language. What experimental psychologists have discovered is that by and large our work-a-day

¹⁵ Rodolfo Llinas, *I of the Vortex: From Neurons to Self*. (Cambridge: MIT Press, 2001).

categories have fuzzy boundaries, and a graded internal structure such that some members are more prototypical than others. Membership is not determined by necessary and sufficient conditions, and categories have a radial structure involving degrees of similarity to the most central members. Moreover these kinds of categories form the fundamental platform for reasoning. ¹⁶This sort of empirical understanding is more akin to exercising a skill than to constructing an argument in discursive form.

Now for inference to the best decision. Notwithstanding Hume's injunction against driving an ought from is, sensible people do often make wise and good decisions about what ought to be done in a social context -- what their children or their nation ought to do or they themselves ought to do. And they do so without invoking normative premises, maxims, rules or what have you. ¹⁷ They judge the situation on the basis of their recognition that one social situation is relevantly similar, perhaps in quite abstract ways, to other social situations whose sequelae are remembered and evaluated. They are using case-based reasoning, not deduction. Case-based reasoning does not require generalizations, normative or otherwise. As Paul Churchland¹⁸ has argued, it probably

¹⁶ See Mark L. Johnson *ibid.*

¹⁷ Paul Churchland "Rules, know-how, and the future of moral cognition. In: *Moral Epistemology Naturalized*, ed. By R. Campbell and B. Hunter, *Canadian Journal of Philosophy* Supplementary Vol. 26, (2000): 291-306. See also Searle's important paper "How derive an ought from an is" *Philosophical Review* Vol. 73 (1964): 43-58. Searle was making a different but related observation. His point was that the existence of certain institutions and standards meant that normal humans often do reason from descriptions about the facts (e.g. I made a promise) to a conclusion about what ought to be the case (I ought to keep my promise). By and large philosophers did not see significant merit in Searle's point. I saw it as the leading edge of a stout and broad-shouldered wedge.

¹⁸ Paul Churchland, *op cit.* 1998.

does depend on fuzzy, radially organized categories whose members occupy positions or trajectories in similarity spaces. (See Figure 1)

The person recognizes a situation as relevantly similar to the cases where courage or kindness or acquiescence or biding one's time was the best strategy. This is cased-based reasoning in decision-making and is very like case-based reasoning in explanation, save that it targets the social domain, rather than the domain of nonsocial phenomena.

Sometimes inference to the best decision calls for action in the very short run, sometimes it calls only for a judgment without immediate action, as when one decides that the USA ought (*morally* ought) to pursue an energy policy that aims for independence on foreign oil. For our purposes, however, the salient point is that these judgments can be understood as inferences to the best decision (as instances of case-based reasoning). As such, they typically make no reference to moral rules or normative generalizations.

What should I do when Adam kicks me? What do I do when I know my close friend is cheating on exams? What do I do if I know my brother murdered someone? These are the questions children ask, and parents, drawing on their own understanding and assessing the complex social situation as best they can, infer to a good decision.

In 1945 Truman decided that the USA would drop an atomic bomb on Japan. Did he *deduce* that the bomb ought to be dropped from facts about the state of affairs? No. Did he use a normative rule plus initial conditions from which the decision followed? No. But he did draw an inference to what he took to be the best decision. The decision was about

what ought to be done, given his beliefs about the stubbornness of the emperor and the widespread Japanese conviction that “better death than dishonor”, his judgment about American casualties in a greatly prolonged war, the likely impact of the bomb’s detonation on Japanese determination, and so forth. On the basis of the same data, others might have come to a different decision, and some would have made the same decision. As with empirical understanding, there are bound to be differences between people regarding the best decision that are not easily resolved. When a collective decision is required despite differing opinions, differences might be safely negotiated.¹⁹ Thus many women who would themselves not wish to have an abortion under any circumstances nevertheless agree that women who feel differently should be allowed the opportunity to abort. Of course sometimes a decision will be negotiated, but once the consequences of the legislation are weighed, the decision may require renegotiation. But many decisions are wise, and by serving the group well, serve the interests of the individual and its capacity to reproduce.²⁰

Does not “best” make this approach circular? No – no more than it does in “inference to the best explanation”. Does it not require too much of us – that our decisions be not only good, but the *best*? No – no more than in inference to the best explanation. “Best” in this

¹⁹ Paul Churchland op cit 1998.

²⁰ As the virtue ethicists have pointed out, there is no evidence that as a matter of fact, Kantians, for example, are more morally upstanding than the general population. In fact, a lot of very ordinary, non-Kantian, people act in morally superlative ways without knowing a thing about the categorical imperative or reflective equilibrium. See Julia Annas, “Being virtuous and doing the right thing.” Presidential Address to the Pacific Division of the APA . In: *Proceedings and Addresses of the American Philosophical Association* Vol. 78, no. 2 (2004): 61-75.

context is just relative to the creature's understanding of the situation, including the social understanding, rooted in neurobiology, of the importance of group flourishing, of the creature's needs, and the desire to avoid suffering. Use the label "case-based reasoning", and the problem disappears anyhow.

An additional point should be made regarding case-based reasoning. Some philosophical discussion of morals seems "academic", in the pejorative sense of the word, because examples presented often lack substantive detail. The lack of detail prevents case-based reasoning from getting traction; too many, or too few, prototypes are available to guide our reasoning. For example, we are asked to consult our "moral intuitions" and render a decision even though the circumstances are grotesquely under described ("when the end of the world is nigh, ought one to punish the guilty anyway?"), or the imagined circumstances are so outlandish that no useful analogies to known experience can be drawn; e.g. suppose you are in a lifeboat and you have to sacrifice Kant or your mother, which should you sacrifice? ²¹

The nature of decision –making, in humans and other animals, is beginning to be understood at both the psychological and the neurobiological levels. The emergence of a social decision depends on background conditions including emotions, motivational factors and social understanding, some of which may be nonconscious and inarticulable.²² It is anchored by social dispositions that are tuned by the reward system

²¹ The so-called "trolley cases" are examples in this vein.

²² See for example, Hsu et al. "Neural systems responding to degrees of uncertainty in human decision-making." *Science* Vol. 310 (2005): 1680-1683; K.D. Vohs, M.L. Meade,

as a result of prior experiences. Given a normal reward system, the young of social species internalize what counts as acceptable behavior; they begin routinely to exhibit it, and they come to expect it from others in the group.²³ We see this in social animals generally – wolves, baboons, vampire bats, Stellar's jays and humans. For example, the developing baboon comes to recognize the requirements of reciprocity in grooming, after having been slapped for failure to reciprocate, just as the child recognizes the paradigmatic requirements of truth-telling after having been punished for lying.

As Aristotle realized, making wise decisions is probably a kind of skill for navigating the social world, and is quite unlike learning the rules of chess or the Ten Commandments. Early moral learning is organized around prototypes and relies on the reward system to make us feel emotional pain in the face of some events (e.g. stealing), and emotional joy in the face of others (e.g. rescuing).²⁴ The child comes to recognize the prototype of being fair, being rude, sharing and cooperating. His understanding is also shaped by the stories of Chicken Little, the Ant and the Grasshopper, and The Little Red Hen. Some understanding may be discursive and rules of thumb may be a provisional guide, in both natural and moral domains: you can't push on a rope; cheaters never prosper. But these are not exceptionless rules and knowing about exceptions is not learning the order in which rules get trumped. Sometimes you can push on a rope (it was wet and is now

M.R. Goode, "The psychological consequences of money." *Science* Vol. 314 (2006): 1154-1165.

²³ Here is the connection I see to Searle's 1964 paper. Once the reward system of the young baboon has been tuned up to the prevailing standard for grooming reciprocity, then the baboon normally acts in accord with the standard. He recognizes a situation wherein another baboon expects grooming for past services, and case-based reasoning takes him to the decision the he ought to reciprocate.

²⁴ See also R.M. Hare *Moral Thinking* (Oxford: Oxford University Press 1981).

deeply frozen) and sometimes cheaters do prosper, though the social cost can be huge. Consequently, much reorganization of understanding occurs as the child matures, just as in the case of his understanding of the natural world. Reflection, mediated by the emulation circuitry²⁵, may occasion significant changes, and some change may occur as part of general, ongoing neuronal consolidation of learning and memory.²⁶ In any case, as Andy Clark wisely argues, case-based reasoning and rule-based (linguaform) reasoning are, at least in modern times and in linguistically competent humans, complementary.²⁷ My point is that case-based reasoning is more fundamental in the cognitive economy.

The role of the brain's reward system in social learning normally engenders respect, or even reverence, for whatever local social institutions happen to exist. Change in those institutions, therefore, may be neither fast nor linear, and may be vigorously resisted even by those who stand to benefit from the change; for example, women who opposed the vote for women; the poor who oppose taxation of the very rich. Despite the power of social inertia, modifications, and sometimes revolutions, do occur, and some of these changes can reasonably be reckoned as moral progress.²⁸

²⁵ Rick Grush "The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences* (2004) 27: 377-442.

²⁶ Some of the consolidation and reorganization likely occurs during sleep. See *The Regulation of Sleep*, A. A. Borbely, O. Hayaishi, T. J. Sejnowski, and J. S. Altman (Eds.), *Human Frontiers Science Program*, Strasbourg, (1996).

²⁷ Andy Clark, "Word and action: Reconciling rules and know-how in moral cognition." In: *Moral Epistemology Naturalized*, ed. By R. Campbell and B. Hunter, *Canadian Journal of Philosophy Supplementary Vol. 26*, (2000): 267-289.

²⁸ Paul Churchland op cit 1998.

A complaint will be that the decisions herein discussed are not genuinely moral, but merely pragmatic. The genuinely moral choice must be made for moral reasons alone. Loftily, the moral purist may insist that one must do the right thing simply and only because it is the right thing, whatever our biology. But real morality is not independent of its relevant domain; namely, getting on in life.

As remarked earlier, our social motivation is rooted in our evolutionary past. Great benefits accrue to organisms living in a social group; individuals share food and other resources, share knowledge of how to hunt and find the essentials of life, and share defense of the group against predators and invaders. We understand reasonably well the conditions permitting social traits to spread throughout a population, and they include the capacity to detect and remember who are the socially dangerous individuals²⁹, the willingness to punish them, and to punish those who will not share the burden of punishment.³⁰

Darwin had the basic story right, when he remarked in *The Descent of Man* (1871) “A tribe including many members who, from possessing in high degree the spirit of patriotism, obedience, courage and sympathy, were always ready to aid one another and to sacrifice themselves for the common good would be victorious over most other tribes; and this would be natural selection.” The very purest of the moral purists have missed the

²⁹ John Batali and Philip Kitcher “Evolution of Altruism in Optional and Compulsory Games” *Journal of Theoretical Biology*, (1995).

³⁰ For a good review of this literature, see *The Origins of Virtue: Human Instinct and the Evolution of Cooperation* by Matt Ridley, (New York: Viking: 1996).

day-to-day nature of real, flesh-and-blood morality. They have certainly missed the significance of the selective advantage generally accruing to individuals who exhibit social behavior, including moral behavior. And I think they have missed the opportunity to understand how morally exemplary behavior can be displayed in, for example, Inuit communities, who lack rule codification but do rely on categories of virtue and case-based reasoning.

Monogamous pair bonding is typical in certain species, such as humans, Canada geese and prairie voles. The behavior exists not because pure reason sees its universal propriety, but because the species evolved so that most individuals have high concentration of receptors for the peptides *oxytocin* and *vasopressin* in limbic structures of the brain.³¹ The limbic pathways connect to the dopamine-mediated reward system (mainly the ventral tegmental area and the nucleus accumbens). Thus, a pair of individuals who copulate come to associate great pleasure with one particular mate, to put the matter simply.

Fundamentally, punishment of cheaters (in the broadest sense) is justified because social traits such as cooperation and sharing cannot spread through a population unless cheaters are punished. Dispositions to punish are likely also to be regulated by neural modulators such as dopamine in the reward system, serotonin in frontal structures and oxytocin in limbic structures. The precise nature of the punishment – shunning, beating, biting or whatever, may, in some species such as humans, be a matter for negotiation and cultural

³¹ L.J. Young, B. Gingrich, M.M. Lim, and T.R. Insel. Cellular mechanisms of social attachment. *Hormonal Behavior* Vol. 40 (2001) 133-139.

standards. Like the importance of punishment in social groups, the higher priority typically given to the welfare of kith and kin over distant strangers, wired up by normal brain development, is unlikely to be significantly changed by any categorical imperative repudiating the morality of such priorities.

A caveat is now in order. Certain social practices may allow the tribe to thrive in one condition, but not another. Lethal outgroup hostility, seen for example in chimpanzees and humans, may be a case in point. According to one model, extinction of competitor out-groups along with high level of in-group sharing, especially of the spoils, gave rise to the distinctive form of altruism seen in human societies.³² As resource availability changes, as the costs of hostile exchanges escalate, as common goals emerge, cooperation may turn out to be a more productive option. Under such conditions, hostile skirmishes acquire moral condemnation and “the melting pot” acquires moral approval. We crack down on urban gangs, and hold out incentives for diversity. Assuming our woodland ape ancestors as well our own human ancestors engaged in outgroup raids, as chimps and several South American tribes still do, can we be confident in moral condemnation of *their* behavior? I see no basis in reality for such a judgment. If, as Samuel Bowles argues, the altruism typical of modern humans plausibly co-evolved with lethal out-group competition, such a judgment will be problematic.

Concluding Remarks:

³² S. Bowles Group competition, reproductive leveling, and the evolution of human altruism.” *Science* Vol. 314 (2006): 1569-1572.

Ultimately, one wants a deeper account of the neurobiological mechanisms of both inference to the best explanation, and inference to the best decision – of case-based reasoning in its many manifestations. My own expectation is that psychology and neuroscience will eventually uncover at least the general principles concerning how neural networks perform these functions and that the two domains (explaining and deciding) will have much in common. My guess, though of course it could be wrong and anyhow is rather vaguely stated here, is that the account will involve neural networks settling into a local minimum in ways that are now quite familiar and well understood. I conclude with this point not because it neatly sews things up, but because I want to emphasize how very much remains to be understood. My main point in this context, however, is that naturalism in ethics should no longer be hobbled by the dictum that you cannot infer an ought from an is. Fine; you cannot deduce an ought from an is. What you can do, however, is come to a decision about what you ought to do without relying on any normative rules or maxims. That is what humans, and undoubtedly other animals, in fact do. From this perspective, many new questions in ethics arise. These questions present philosophers with a unique opportunity to collaborate with scientists on matters of great social importance.

Figure1. Caption:

A schematic diagram characterizing the parameter space for that subset of social categories with ethical relevance. The aim of the diagram is purely conceptual; namely, to augment the text by showing what is meant by a parameter space, how categories cluster, and how similarity can be understood in terms of distance in a parameter space. It relies on principles of the parameter-space approach to representation in sensory systems, such as taste, vision and audition, or to representation of more complex categories such as faces and furniture. In sensory systems, there is a lot of evidence from neurophysiology supporting the hypothesis of parameter-space representation. The morally blameworthy categories are black, and cluster on the far side of the main divide, and the morally praiseworthy actions cluster on the near side of the divide. Adapted from Paul Churchland, *op. cit.* p. 87.