

- Stoerig, P. and Cowey, A. (1997) Blindsight in man and monkey. *Brain*, 120: 535-559.
- Treisman, A. (1988) Features and objects: The Fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, 40: 201-237.
- Treisman, A. and Gelade, G. (1980) A feature integration theory of attention. *Cognitive Psychology*, 12: 97-136.
- Treisman, A. and Gormican, S. (1988) Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95: 15-48.
- Trentman, J. (1970) Ockham on mental. *Mind*, 79: 586-590.
- Ungerleider, L.G. and Mishkin, M. (1982) Two cortical visual systems. In D.J. Ingle, M.A. Goodale, and R.J.W. Mansfield (Eds.) *Analysis of Visual Behavior*. Cambridge, MA: M.I.T. Press.
- Vanni, S., Revonsuo, A. and Hari, R. (1997) Modulation of the parieto-occipital alpha-rhythm during object-detection. *Journal of Neuroscience*, 17 (18): 7141-7147.
- Weiskrantz, L. (1986) *Blindsight*. Oxford: Clarendon Press.
- Wilkes, K. V. (1984) Is consciousness important? *British Journal of the Philosophy of Science*, 35: 223-243.

MITCHELL R. HANEY

## Dynamical Cognition, Soft Laws, and Moral Theorizing

There are two aims of this paper: First, I will offer a critique the current discussion of connectionism's implications for moral theorizing. Second, I will outline a position concerning the proper structure of moral theory that is suggested but not fully developed in Terry Horgan and John Tienson's *Connectionism and the Philosophy of Psychology* (1996) A 'Soft Principles' structure of moral theorizing is a position that has largely been dismissed by ethical theorists, but it is a position that through an analogy with Horgan and Tienson's account of human cognition and psychological theorizing may receive renewed attention.

Keywords: connectionism, ethical theory, particularism, mind, morals

In ethics ... there are no hard rules.  
There are only *ceteris paribus* generalizations.  
*Terence Horgan and John Tienson*

### 1 Connectionism and Moral Theory: The Current Discussion

For those philosophers wedded to methodological naturalism, any progress concerning our understanding of the natural world provides new data and conceptual tools for other arenas of philosophical inquiry. In recent years, the progress in our understanding of the human mind via connectionist models of human cognition has begun to have conceptual repercussions in the arena of moral theorizing. A few philosophers have sculpted a landscape of meta-ethical positions concerning the structure of moral theorizing based upon results and conceptual equipment gleaned from connectionist models of human cognition.

It is often claimed of modern moral theorizing that an adequate moral theory must be a tightly hierarchical set of moral principles and rules that capture morally relevant features of our environment, and that such features provide neatly defined sets of necessary and sufficient conditions for the application of moral concepts. (E.g. maximizing utils provided the necessary and

sufficient conditions of moral rightness in Bentham's rendition of utilitarianism.) Mark Johnson has called this presupposition of much modern moral theorizing, "Moral Law Folk Theory." (Johnson 1993, 4) Johnson states that moral theorists have viewed moral judgment as, "the judging of particular cases as falling under a particular moral concept, and thereby being governed by a specific moral rule." (207) There is much to the claim that moral theorists have, at least, implied that rational moral judgment is, properly speaking, an algorithmic process. However, if one is a methodological naturalist and one is persuaded by the likelihood that connectionism offers a true account of cognitive processing, then things look bleak for moral law folk theory.

There is a general, sometimes sketchy, line of reasoning against principle and rule based moral theorizing, i.e. theorizing guided by the assumptions of moral law folk theory, that can be gleaned from connectionist writings on moral theorizing. The standard argument appears to be the following: Moral theories primarily aim to provide models of rational moral deliberation. Modern moral theorists have offered algorithmic or subsumptive decision-procedures (e.g. the utilitarian calculus, universalizability, etc.). And, such algorithmic decision-procedures require that moral features (good/bad, right/wrong, etc.) be captured in moral principles or rules that contain the necessary and sufficient conditions of their application (e.g. an act is right iff the act intention can be willed to be a universal law, or an agent is good iff she intends to maximize happiness for all affected by her actions). However, connectionist models of cognition — if correct — conclude that the bulk of human reasoning and inference is non-algorithmic in nature. First, connectionists contend that belief formations and transitions are predominantly based in pattern completion and extension rather than deductions from subsuming cases under principles or rules. Second, connectionists argue that the concepts employed within human reasoning are chiefly exemplary or prototypical in nature; hence, most of our concepts fail to be analyzable into necessary and sufficient conditions. As a result of these findings, connectionists tend to conclude that the modern moral theorist's understanding of moral judgment is truly misguided, because it rides upon a false understanding of the nature of human reasoning and concepts. Given that the connectionist understanding of human cognition mitigates the plausibility of moral law folk theory being true, some connectionists have aimed to offer new models for the structure of moral theorizing informed by their model of human cognition.

There are two main genera of connectionist accounts of moral theorizing. The first captures those models that take a generally eliminativist attitude toward moral principles and rules. The second captures those models that take a generally non-eliminativist attitude toward moral principles and rules. Within the genus of eliminativists there are both global and local species. A global eliminativist, which the views of Paul Churchland tend to approach, holds that there is little or no role for moral principles and rules. According to Churchland, moral deliberation is specifically a matter of exemplar driven cognition. Hence, Churchland claims that moral theory ought to be structured more along the lines of Aristotle's virtue ethics than Kant's system of categorical imperatives. (Churchland, 1989) The local eliminativist, which the work of Hubert and Stuart Dreyfus is representative, argues that there is some role for moral principles and rules but they are, in the end, eliminable. For the Dreyfuses, agents must rely upon moral principles until they master morality in such a way that exemplar driven cognition takes over. (Dreyfus and Dreyfus, 1990) In short, the Dreyfuses argue that moral principles are merely heuristic devices to be replaced by more sensitive cognitive processing as agents morally develop. In sum, the eliminativists find that, at least, mature moral deliberation is as connectionists suggest of most developed reasoning—it is exemplar driven, and that moral principles and rules, at most, play a heuristic role in moral development.

There are two species of non-eliminativism with regard to connectionism's implications for moral theorizing. These species are: Local Compatibalism and Global Compatibalism. The first specie is an occupied position on the terrain, and the second I shall develop later in this paper.

The local compatibalist, of which Andy Clark's view is most representative, argues that moral principles are ineliminable but — on the whole — play a much different role in moral reasoning than has traditionally been thought. Clark shares with the eliminativists that moral principles and rules are intrapersonal heuristic devices used in the training as well as monitoring of individual cognitive systems, but he argues that moral principles are ineliminable from interpersonal problem-solving situations. Clark states that in problems to be solved by groups rather than individuals moral principles and maxims "lay out some rough guides and signposts that constrain the space to be explored in the search for a cooperative solution." (Clark 1996, 122) As a result, Clark argues that moral principles are ineliminable but their role is primarily not for individual decision-making; instead, the use of moral principles (as linguistic expres-

sions) is to massage and shape the direction of social cooperation. The global compatibalist, as I shall argue below, holds that moral rules are ineliminable from moral theorizing, and that they play roles tantamount to those assigned to moral principles within traditional moral theorizing. However, before I move on, I have a bone to pick with the current discussion of connectionism's implications for the structure of moral theorizing.

## 2 Problems for the Current Discussion

The standard argument sketched above, I shall suggest, fails to undermine principle based moral theory in a way that calls for a radical revision of the entire structure of moral theorizing. It fails to do so for a number of reasons, but all are largely related to the fact that assumptions encapsulated in Johnson's account of Moral Law Folk Theory are easily challengeable by traditional moral theorists.

First, Johnson, as well as many connectionists, presuppose that ethical theory demands that moral concepts have necessary and sufficient conditions analyses in order to yield moral principles. This is an inaccurate assessment of some leading modern moral theories. For instance, an intuitionist like G.E. Moore totally denies that the good can be defined or illuminated by a set of natural features that elucidate the necessary and sufficient conditions of its application. (Moore 1903) Likewise, a 'constructivist' like Hilary Putnam denies that there is any such analysis of the Good or what he calls, "human flourishing". (Putnam 1981, 127ff) W.D. Ross goes one step further and claims that both the Right and the Good fail to be subject to such an elucidation. (Ross 1930) And, the Kantian, Alan Donagan states:

The concept of respecting a human being as a rational creature is not usefully definable for our purposes. Thus to define it as treating a human being, by virtue of his rationality, as an end in itself, while perhaps clarifying, does not furnish us with a useful substituent. Yet, it does not follow that the process of deriving specific precepts from the fundamental principle is arbitrary and unreasoned. (Donagan 1977, 67)

These references show that it is not the case that modern moral theorists have unequivocally held that moral features or even morally relevant features (e.g.

respect for persons) can be decomposed into sets of necessary and sufficient conditions. In fact, these references suggest that moral theorists need not hold the moral law folk theory view and that they may hold that what we have in the realm of the moral are prototypical instances that come to represent our moral concepts. For instance, Ross claims that we don't have a definition of the good, but we have things that are — one might say — prototypically good. For Ross, the things that are prototypically good include: virtue, pleasure, allocation of pleasure to the virtuous, and knowledge. (For which I am sure there maybe prototypes of each of these, rather than definitions.) The thing to note here is that despite the fact that each of the above theorists admit that central moral concepts are indefinable and may allow for prototype or exemplar based illuminations of these concepts, none of these theorists abandon principle or rule based structures of moral theorizing. Moore embraces the principle of ideal utility to determine right actions. Ross provides his famous list of six *prima facie* duties. And Donagan argues that the second formulation of Kant's categorical imperative is both morally fundamental, as well as useful for determining the morality of actions. Hence, it is not necessary for moral theorists to abandon principled accounts of morality, if the concepts they reflect are not subject to necessary and sufficient conditions analyses.

Second, there is a distinction that has long been familiar within the arena of ethics that takes the bite out of what I have presented as the connectionist's standard argument against modern moral theorizing. Ethicists distinguish between principles as the basis of a moral decision-procedure and principles as capturing moral criteria. (Bales, 1970) It is true that ethicists have often thought that fundamental moral principles are both the criteria of morality and the basis of rational moral deliberation, but — conceptually — they are distinct. David Brink states of moral principles that they "can be construed as *standards* or *criteria of rightness* or as *decision procedures*. A standard or criterion of rightness explains what makes an action or motive right or justified; a decision procedure provides a method of deliberation." (Brink 1989, 216) The connectionist's argument against modern moral theorizing, as I have elucidated it, appears to rest on the mistake that ethical theories are only in the business of providing moral decision-procedures. Insofar as ethical theories also aim to understand what makes right acts right and good character traits good, moral principles may capture the criteria or standards of right actions, but be useless as a means of moral deliberation, i.e. they need not be cognizable. As a result,

simply arguing that the way in which we judge the morality of actions and character traits is not capturable by exceptionless moral principles fails to undermine principles as capturing the criteria of morality. It may be argued that moral deliberation is much more effective in meeting a moral criterion, e.g. the principle of utility, if it utilizes a rich set of moral exemplars rather than principles or rules.<sup>1</sup>

To illustrate the above point let us consider the following: Suppose we are utilitarians and let us grant that moral decision-making cannot be strictly or even largely algorithmic in nature in virtue of the fact that our reasoning capacities are actually exemplar or prototype driven. Does this fact undermine the utilitarian claim that our decisions or actions may still be right/wrong in virtue of whether or not they in fact maximize happiness? It seems that the answer is, 'No'. It does not matter how we make moral judgments on the viewpoint that the principle of utility captures a moral criterion, it only matters whether or not the outcomes of our deliberations in fact meet the criterion. If this is so, then the connectionist's standard argument against principled based moral theory fails to undermine the entirety of the modern structure in the way that they seem to think it does. If their argument is correct (and I believe it is), then it only calls for moral theorists to abandon the aim of providing algorithmic decision-procedures, but need not motivate theorists to abandon principle structured accounts of moral criteria. As a consequence, any of the connectionist models of moral deliberation outlined above are completely compatible with a view of principled based moral theories as providing an account of that in virtue of which acts, persons, etc. have the moral properties they have. In the next section, I shall suggest an additional argument for favoring a more global rejection of principled-based models of moral theorizing.

### 3 Dynamical Cognition and Moral Particularism: An Analogy

As I have argued above, the connectionist critique of moral theory via an attack on principle based moral decision-procedures fails to necessitate the wholesale revision of moral theorizing. However, eliminativists and non-eliminativists

<sup>1</sup> This is not a new point within moral theorizing. Mill, Sidgwick, and Kant all consider the possibility that their criteria of morality fail to be useful for ordinary human judgment, and all propose things like rules of thumb or imperfect duties to cover for the possible intractability of their moral principles as decision-procedures.

alike often suggest that moral theorizing must be radically revised on the basis of connectionist accounts of moral deliberation. In opposition to wholesale revision, I concluded that moral theorists may only be asked to abandon algorithmic models of moral deliberation, and that moral principles viewed as moral criteria are left untouched by what appears to be the standard connectionist argument. There is, however, I think an argument related to connectionist understandings of human cognition that does suggest a global re-orientation of moral theorizing. Such a tack toward moral theorizing is suggested, but never developed by Horgan and Tienson when they state that:

(I)t is entirely possible that normative standards are like competent human cognition, in this respect — i.e., that normative standards are too complex, too subtle, and too sophisticated to be formulated as exceptionless general principles. (Horgan and Tienson, 143)

An additional argument against modern moral theorizing that is symbiotic with connectionist frameworks would be an argument by analogy. The main claim of this argument is as follows: If moral phenomena display structural complexities analogous to human cognition, and connectionism has an adequate model of human cognition, then we ought to be able to illuminate the structure of morality using a similar model.<sup>2</sup>

Connectionists claim that belief formations and transitions are a matter of pattern completion and extension, and that the patterns are a matter of having concepts that are exemplary or prototypical in structure and that serve as grounds for reasoning and inference. What makes prototypes interesting is that they fail to have necessary and sufficient conditions of application. The interest in prototypes is — in part — due to the fact that they allow that there can be cases where certain properties of a prototype are present in the environment but the concept represented in the prototype is not activated. As a result, human cognition is much more sensitive to the subtleties and intricacies of the environment than exceptionless generalizations could ever capture. For instance, consider Churchland's famous Mine — Rock Device. (Churchland 1988, 157-162)

<sup>2</sup> I would like to note that though I shall be employing the lexicon of moral realism, e.g. I shall talk of "moral properties", etc., I take it that the metaphysics of morality is orthogonal to issues regarding the structure of moral theorizing. For instance, I could also argue in an irrealist way that our deep moral norms or norms reflecting our emotive responses to situations are, at root, never exceptionless.

The responses of the mine detector when it bounces off a mine as opposed to when it bounces off a rock are so similar that there is no known rule that could state the sufficient differences between a mine and a rock. As a result, a mine response and a rock response are very similar indeed; thus, each prototype will overlap in many ways. However, a connectionist network is, after sufficient training, able to detect the subtle differences between mine and rock responses. What this points to is that some subtle difference in the environment can lead to a different response by the mine-rock device, and it is the subtle differences that make all the difference between running into rocks and running into mines. Looking to the recent theory called moral particularism, it is argued that morality has a similar contextually sensitive structure.

Moral particularists claim that there are no universally relevant non-moral features. (Dancy, 1993; McNaughton, 1988) What they mean by this is that there are no non-moral features that when present are always morally relevant and always relevant in the same way. For instance, moral particularists would claim contra hedonism that pleasure is not always morally relevant or always a right-making feature when instantiated or instantiable. Instead, the particularists argue that there are sets of features that are prototypically relevant, but all can be *silenced* or *reversed* given the presence of other non-moral features in the circumstances.<sup>3</sup> As a result, pleasure may be prototypically relevant feature, but it may be silenced or irrelevant when it is coupled with another feature, e.g. act of denying freedom of movement. (The fact that someone experiences pleasure while physically restraining a non-compliant, innocent individual apparently doesn't mitigate the immorality of the act.) To use another simple example: Intentionally telling a falsehood is typically morally relevant and typically a wrong-making feature. However, in cases of white lies the intentional telling of a falsehood is typically thought either to be silenced, i.e. it is irrelevant, or overridden, i.e. relevant but not paramount, because of some other feature of the circumstance is more morally salient. In the case of white lies either avoiding harm or producing some good is thought to be so paramount in the circumstances that it either silences or overrides the typical moral relevance of intentionally telling falsehoods. In specific cases such as

<sup>3</sup> I am employing a moderate form of moral particularism that allows that there may be rules of moral relevance that pick out the tendencies for some non-moral features to be morally relevant. More radical forms of moral particularism argue that morally relevant features cannot be determined — even defeasibly — apart from particular cases. See Dancy, 1993.

these, we have instances of a more general structure. To put it in connectionist terminology, these are cases in which prototypically relevant features are either not morally activated or have little moral weight due to differences — however small — in other non-moral features of the circumstances which counter-act or weigh against them. As a result, the overall morality of an act or a character trait is determined by the settled weights of all the activated features in the circumstances. If we couple such simple examples, and many others like it, with the history of debates between substantive ethical theories, then we get an even more evidence for the claim that there may be some non-moral features that are prototypically morally relevant but never universally relevant.

Much of the history ethical theorizing has centered on the use of counter-examples to undermine the credibility of opposing ethical theories. Counter-examples typically exhibit the following structure: They use an opposing ethical theory's principle(s) and apply it to a case where the resulting conclusion is counter-intuitive. The intended result is that we ought to think that the feature captured in the moral principle is not relevant in the case; hence, it is not universally relevant in the way that a given theory claims. For instance, H.J. McCloskey presented a hypothetical case meant to undermine the claim that maximizing utility is not a universally relevant non-moral feature. (McCloskey 1959) He asks his reader to consider a case where a utilitarian is visiting an area suffering from racial strife. The utilitarian is in this area when a male of the minority group rapes a female of the majority group. Race riots ensue, and the majority group with help from the local authorities maim and kill members of the minority group. Since the utilitarian was in the vicinity of the crime, her testimony against any male member of the minority group would bring about his conviction, and it would also bring an end to the riots. McCloskey claims that the good act-utilitarian must conclude that she has a moral duty to bear false witness against an innocent person, because that would maximize utility for all affected. The specific point of McCloskey's example is that maximizing utility fails to account for considerations of justice, e.g. that bearing false witness is wrong. However, the more general claim is that the feature of maximizing utility is not always relevant or, at least, overriding, and that is the central tenet of utilitarianism. It has been strongly suggested by the history of ethics that similar counter-examples can be generated against all principles forwarded by the variety of substantive moral theories that are and have been defended. As a result this history, it is plausible that there are no features that are universally

relevant, and that morality is a much more complex phenomena than has traditionally been supposed.

The history of debate between substantive ethical theories coupled with the moral particularists's specific claim that though there may be prototypically morally relevant features but no universally relevant ones suggests that morality has a structure analogous to cognition as described by connectionism. That is, morality may have a kind of open-ended structure in order that always allows for non-prototypical cases, and that any theory adequate to the phenomena must have the flexibility to encompass nuances of particular circumstances. To put this in the terms Horgan and Tienson employ, it may be very likely that morality is dynamical in nature, i.e. that the morality of acts, character traits etc. emerge from complex networks of prototypically morally relevant features which may take on different moral weights under different particular configurations. If it is the case that morality and cognition are structurally analogous insofar as they are both dynamical in nature, then we may consider the claim that moral theory take on a structure similar to connectionist models. If the foregoing claim is plausible, then — at the very least — morality cannot be captured via exceptionless moral principles.

The above argument by analogy aims to achieve what connectionists hoped for — it is an argument that suggests that we need to revise the structure of all moral theorizing. If moral phenomena — in the broadest sense — are too complex and subtle to be captured by exceptionless moral principles or rules, then not only are exceptionless principled accounts of moral deliberation wrong-headed, but so are exceptionless principled accounts of moral criteria. With this argument, if it is in fact cogent, we are in a position to accept some connectionist model of moral theorizing whether it be eliminativist or non-eliminativist. However, I think that there are some reasons for extending Horgan and Tienson's soft laws account of psychological theorizing to moral theorizing. Such a view, I shall suggest, fills in the neglected option under non-eliminativism I am calling Global Compatibalism.

#### 4 Soft Laws and Moral Theorizing

If moral phenomena are structurally analogous to cognition, and Horgan and Tienson are correct in the fact that certain cognitive forces lead to defeasible causal tendencies, then it may well be the case that certain non-moral features

have defeasible relevance relations with regard to the deontic status of actions and the morality of character traits. In fact, I shall argue that the history of ethical theory makes plausible the claim that there are certain features that are defeasibly relevant. But, first consider the basics of Horgan and Tienson's account of psychological theorizing.

Horgan and Tienson argue that inspite cognition's being dynamical in nature there are defeasible, psychological causal tendencies. The natural expression of any defeasible causal tendencies within the cognitive realm, according to Horgan and Tienson, is through *ceteris paribus* generalizations, and that there are some *ceteris paribus* generalizations that "by virtue of their scope and generality" deserve to be considered laws, i.e. soft laws. (Horgan and Tienson 1996, 107) In addition, Horgan and Tienson claim, "Since the basic causal tendencies of cognitive systems come in systematic arrays, cognitive systems are (fundamentally) characterizable by *parameterized ceteris paribus* generalizations." (115) Finally, due to the open-endedness of human cognitive capacities, as well as the frame-problem, Horgan and Tienson conclude that there is no way to state the exceptions that would capture all possible cases where *ceteris* is not *paribus* in the realm of cognition. (31ff) I shall suppose that Horgan and Tienson are correct in their account of the appropriate model for theorizing about the nature of cognition and laws of cognition. Hence, it remains to be seen whether or not morality may not be subject to a similar structural model.

Consider the history of debate between substantive ethical theories once again. It seems that despite the fact that counter-examples in ethics do lead to counter-intuitive conclusions they never seem to undermine our confidence in the fact that a feature captured by a principle in question tends to be morally relevant. For instance, despite the fact that McCloskey's counter-example yields a biting challenge to the universal relevance of maximizing utility, we remain largely convinced that maximizing utility is still an important morally relevant feature. The same holds true for many other features captured by moral principles and rules, e.g. respect for persons, intentional telling of falsehoods, benevolence, etc. Everyone of these features has counter-examples suggesting that they are not always morally relevant, but in spite of that, we remain convinced that they tend to be relevant and tend to be relevant in certain ways. What I suggest is that, analogous to cognitive forces having defeasible causal tendencies, certain non-moral features have defeasible tendencies with regard to their

moral relevance. This model of the structure of moral theory is not entirely alien in the history of ethics.

W.D. Ross defended the view that moral normativity cannot be reduced to or captured by a single, universally relevant feature. Ross states, after arguing that neither monistic consequentialism nor monistic deontology are free from counter-examples, that "... it is more important that our theory fit the facts than it be simple ...". As a result, Ross abandons monism for pluralism in ethics, because, he seemed to believe, that examples both for and against individual features show both the underlying truth of the typical moral relevance of certain features, as well as the fact that they are not always morally overriding. In addition, he proposes that due to the fact that the set of features upon which the moral properties of actions typically supervene are defeasible that they are to be captured by rules that contain ineliminable *ceteris paribus* clauses. For instance, fidelity (truth-telling and the keeping of promises), according to Ross, is *ceteris paribus* morally relevant and a right-making feature. The same goes for Ross's five other morally relevant features of right action. As a result, pluralism and non-exceptionless rules made up the structure of Ross's system of *prima facie* duties. And Ross states, "For the comparative stringency of these *prima facie* obligations no general rules can, so far as I can see, be laid down." Hence, according to Ross, the exceptions to the *prima facie* duties, i.e. when *ceteris* is not *paribus*, cannot be formulated or determined a priori. Hence, *ceteris paribus* moral generalizations are fundamental to the description of morality. As a result of all this, it appears that Ross's account of morality and moral theorizing is analogous to Horgan and Tienson's account of cognition and psychological theorizing.

Now, Ross's *prima facie* duties represent one set of *ceteris paribus* moral generalizations with regard to right actions, but one could expand the list to include others, e.g. features related to good character traits. In addition, one could expand or contract the entire list based upon moral experience, moral discourse, and the armchair discoveries of moral theorists. For instance, in recent decades moral theorists have brought to the fore such neglected morally relevant features as care and trust. Here I will not argue for what content ought to be included in the set of soft moral generalizations. Instead, it remains to be argued how the view outlined above fills out the option of connectionist moral theorizing that I dubbed Global Compatibalism.

If I am correct that a Rossian style ethical theory is structurally analogous to Horgan and Tienson's account of psychological theorizing, then we attain a renewed, connectionist symbiotic approach to moral theorizing that is a more robust than any of the present forms of eliminativism or non-eliminativism with regard to the role of moral generalizations. Under Horgan and Tienson's account of psychological theorizing, soft laws do sufficiently describe psychological causation for the purposes of explaining cognitive processing. And if the analogy I have been articulating holds, then Rossian-like *ceteris paribus* moral generalization may actually capture the *de facto* tendencies of some set of non-moral features to be morally relevant (even though they are not universally relevant). As a result, such generalizations may be more robust than either being merely heuristic devices within moral development or linguistic expressions useful for attaining and maintaining social cooperation. Rossian *ceteris paribus* moral generalizations, let us call them 'Soft Moral Principles', may actually fulfill the role of moral criteria in the face of the complexities of moral phenomena that exceptionless moral principles fail to capture. In the way that soft laws capture defeasible causal forces within the realm of the cognitive, according to Horgan and Tienson, and, thus, constitute the backbone of psychological theorizing, it may well be that soft moral principles capture the defeasible relevance relations of morality and, thus, constitute the backbone of ethical theorizing. As a result, on analogy with Horgan and Tienson's approach to psychological theorizing, we may attain a model of moral theorizing that is sensitive to connectionist data and conceptual tools, but maintains a more robust and global place for moral rules than any other connectionist-inspired position to date.<sup>4</sup>

## References

- Bales, R.E. (1971). "Act Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?". *American Philosophical Quarterly* 8: 257-265.
- Bentham, J. (1948). *The Principles of Moral and Legislation*. New York, NY: Hafner Publishing Company.

<sup>4</sup> I want to thank Terry Horgan and John Tienson for their philosophical guidance and inspiration over the past years. I also want to thank Matjaž Potrč, Olga Markič, the Department of Philosophy at the University of Ljubljana, as well as the conference participants for a successful and informative conference. Last, but not least, I would like to thank the Ethics Center at the University of South Florida and the following friends for their helpful comments and support: Peggy DesAutels, Bob Ennis, Robin Fiore, Herman Stark, and Margaret Walker.

- Brink, D. (1985). *Moral Realism and the Foundations of Ethics*. New York, NY: Cambridge University Press.
- Churchland, P. (1988) *Matter and Consciousness*. Revised edition. Cambridge, Mass: MIT Press.
- (1989). "Moral Facts and Moral Knowledge". *A Neurocomputational Perspective*. Cambridge: MIT Press.
- Clark, A. (1996) "Connectionism, Moral Cognition, and Collaborative Problem Solving". *Mind and Morals: Essays on Cognitive Science and Ethics*. Larry May, Marilyn Friedman, and Andy Clark, eds. Cambridge, Mass.: MIT Press.
- Dancy, J. (1983). "Ethical Particularism and Morally Relevant Properties". *Mind* XCII: 530-547.
- (1993). *Moral Reasons*. Cambridge: Basil Blackwell.
- Donagan, A. (1977). *The Theory of Morality*. Chicago: The University of Chicago Press.
- Dreyfus, H. and S. Dreyfus, (1990). "What is Morality? A Phenomenology of the Development of Ethical Expertise." in *Universalism vs. Communitarianism: Contemporary Debates in Ethics*. Cambridge, Mass: MIT Press.
- Horgan, T and Tienson, T. (1996). *Connectionism and the Philosophy of Psychology*. Cambridge, Mass: MIT Press.
- Johnson, M. (1993). *Moral Imagination: Implications of Cognitive Science for Ethics*. Chicago: University of Chicago Press.
- Kant, I. (1964). *Groundwork of the Metaphysics of Morals*. H.J. Paton (trans.), New York: Harper & Row.
- (1991). *The Metaphysics of Morals*. M. Gregor (trans.), New York: Cambridge University Press.
- McCloskey, H.J. (1965) "A Non-Utilitarian Approach to Punishment," *Inquiry* 8: 239-255.
- McNaughton, D. (1988). *Moral Vision: An Introduction to Ethics*. New York: Basil Blackwell.
- Mill, J.S. (1979). *Utilitarianism*. Indianapolis, IN: Hackett Publishing.
- Moore, G.E. (1913) *Principia Ethica*. Indianapolis, IN. Hackett Publishing.
- Putnam, H. (1981). *Reason, Truth, and History*. New York, NY: Cambridge University Press.
- Ross, W.D. (1988). *The Right and the Good*. Indianapolis: Hackett Publishing.
- Sidgwick, H. (1981). *The Methods of Ethics*. (7th ed), Indianapolis, IN: Hackett Publishing.

HERMAN E. STARK

## What the Dynamical Cognitive Scientist Said to the Epistemologist

I explore what Terry Horgan and John Tienson's Dynamical Cognition framework suggests for the theory of rationality. I enumerate, in other words, various ways in which the "DC" framework constrains and opens up the attempt to understand what it is to be rational. Part of my discussion draws on what Horgan, in collaboration with David Henderson, has written on the topic. The main result of my efforts is the development of a list of DC suggestions for epistemology that are conjunctively formidable and furthermore illustrative, both generally and specifically, of why epistemologists cannot afford to ignore research in cognitive science.

Keywords: epistemology, connectionism, philosophy of psychology, the implications of connectionism for epistemology, dynamical cognition and being justified

In middle and late twentieth century philosophy of science philosophers have chided philosophers for having had the audacity to tell scientists how science should proceed. The new philosophy of science instead suggested that philosophers stop running ahead of science and rather accompany science as the latter made its way to wherever it is going. But this peaceful and pace-respecting relation has perhaps lulled the philosophers into a slumber, for now some of the sciences have crawled into the lead and are telling philosophers how philosophy should proceed. In this essay I consider one example of this crafty revenge. I look at what Terry Horgan and John Tienson propose, in their *Connectionism and the Philosophy of Psychology*, as the proper framework for understanding human cognition (the Dynamical Cognition or "DC" framework), and then observe how both they (in that same book), David Henderson and Terry Horgan (in "Iceberg Epistemology"), and I myself (in this essay) take this DC framework to be variously constraining and indeed variously shaping what counts as legitimate epistemology. But since epistemology is a big field and this is a small essay, the main focus herein will be on the analysis of "being justified," and thus the overall purpose of the essay can be described as showing how the DC



