

Notes on Practical Reasoning

Gilbert Harman

June 13, 2007

1 Reasoning by oneself and reasoning with others

In these notes, I will use the word “reasoning” to refer to something people do. The general category includes both internal reasoning, reasoning things out by oneself—inference and deliberation—and external reasoning with others—arguing, discussing and negotiating.

Accounts of reasoning of either sort can be either descriptive psychological theories that attempt to characterize something about how people actually infer, deliberate, and argue, or normative theories that attempt to say something about how people ought to infer, deliberate and argue, or to characterize certain aspects of good or bad inference, deliberation, or argument.

The word “reasoning” can also be used to refer to a product of the process of reasoning. Something written down, for example. I don’t want to confuse the process with the product, so I here use “reasoning” just for the process.

The product of reasoning might be a linear sequence of sentences or propositions, perhaps in the form of a linear argument. It does not follow that the process of reasoning is linear. It almost certainly involves highly parallel processes of constraint satisfaction.

2 Logic, probability, and decision theory

I need at the start to say something about the relevance of logic, probability, and decision theory to reasoning by oneself—inference and deliberation—and to reasoning with others. The matter is complicated because the terms “logic”, “theory of probability”, and “decision theory” are used sometimes to refer to formal mathematical theories of implication and consistency, sometimes to refer to theories of method or methodologies, and sometimes to refer to a mixture of theories of these two sorts.

On the formal mathematical side, there is formal or mathematical logic, the mathematical theory of probability, and mathematical formulations of decision theory in terms of maximizing expected

utility. It is an obvious point, which is nevertheless often missed, that these formal theories are neither descriptive theories about what people do, nor normative theories about what people ought to do. So they are not theories of reasoning in the sense in which I am using the term “reasoning.”

Although accounts of formal logic (e.g., Goldfarb 2003) sometimes refer to “valid arguments” or examples of “reasoning” with steps that are supposed to be in accord with certain “rules of inference,” the terms “reasoning” and “argument” are then being used to refer to certain abstract structures of propositions and not for something that people do, not for any concrete process of inference or deliberation one engages in by oneself or any discussion among two or more people. The rules in question have neither a psychological, nor a social, nor a normative subject matter. They are rules of implication or rules that have to be satisfied for a structure to count as a valid formal argument, not rules of inference in the sense in which I am here using the term “inference”.

Properly stated, the logical rule of *modus ponens* says that a conditional and its antecedent jointly imply the consequent of the conditional. The rule does not say that, if one believes or otherwise accepts the conditional and its antecedent, one must or may also believe or accept the consequent. The rule says nothing about beliefs and nothing about what may or may not be asserted in an argument in our sense.

There may be corresponding principles about what people do or can or should rationally believe or assert, but such principles would go beyond anything in formal logic. Indeed, it is nontrivial to find corresponding principles that are at all plausible (Harman 1986, Chapter Two). It is certainly not the case that, whenever one believes a conditional and also believes its antecedent, one must or may rationally believe its consequent. It may be that one also already believes the negation of the consequent and should then either stop believing the conditional or stop believing its antecedent.

A further point is that inference takes time and uses limited resources. Given that any given set of beliefs has infinitely many logical consequences, it is simply not true that one rationally should waste time and resources cluttering one’s mind with logical implications of what one believes.

Similar remarks apply to consistency and coherence. Formal logics, probability theories, and decision theories characterize consistency of propositions and coherence of assignments of probability and utility. Such formal theories do not say anything about what combinations of propositions people should or should not assert or believe or about what assignments of probability and utility they should accept. There may be corresponding principles connecting consistency and coherence with what people should rationally not believe or assert, but again those principles go beyond anything in formal logic, probability theory, and decision theory and again it is nontrivial to find such principles that are at all plausible.

Given limited resources it is not normally rational to devote significant resources to the computationally intractable task of checking one’s beliefs and probability assignments for consistency and coherence. Furthermore, having discovered inconsistency or incoherence in one’s beliefs and assignments, it is not always intelligent to drop everything else one is doing to try to figure out the best way to eliminate it. The question of what to do after having discovered inconsistency or incoherence is a methodological issue that can be addressed only by a normative theory of reasoning. The answer is not automatically provided by formal logic, probability theory, and decision theory.

As mentioned earlier, the terms “logic”, “probability theory”, and “decision theory” can be used not only for purely formal theories but also for methodological accounts of how such formal theories might be relevant to rational reasoning and argument (Mill 1846, Dewey 1938). These methodological proposals are additions to the purely formal theories and do not follow directly from them.

3 Theoretical, Practical, and Moral

We can distinguish “theoretical” reasoning or inference from “practical” reasoning or deliberation. Theoretical inference may result in new beliefs, or changes in old beliefs, whereas practical deliberation may result in new decisions, intentions, and plans or changes in old decisions, intentions, and plans. (It is of course possible that the process does not lead to any changes of these sorts.)

External theoretical discussion with others is typically expressed in declarative sentences, like “Albert went to the late show at the Garden Theater,” whereas practical discussion with others often makes use of imperatives, like “Let’s go to the late show at the Garden!”

Moral reasoning can be either theoretical or practical—theoretical if the issue is what (to believe about what) someone ought to do, or what is good or bad, right or wrong, just or unjust; and practical if the issue is what to do when moral considerations are or might be relevant.

4 Internal Reasoning

Internal reasoning—reasoning something out by oneself, inference and personal deliberation—is not the same thing as external reasoning—bargaining, negotiation, argument, justification (to others), explanation (to others), and other sorts of discussion among various people.

Internal inferences and deliberations are processes that typically lead to (relatively small) changes in one’s attitudes (beliefs, intentions, desires, etc.) by addition and subtraction. They are not the only such processes. One can forget things and can also suffer from illness and injuries leading to more drastic changes that are not instances of reasoning.

It is unclear how to distinguish processes of internal reasoning from such other processes. For one thing, there appear to be rational ways of acquiring beliefs as “direct” responses to sensation and perception. Are these instances of reasoning? The matter is complicated because unconscious “computation” may occur in such cases and it is difficult to distinguish such computation from unconscious reasoning. A similar point holds concerning “intuitions.”

5 “Infer,” “inference,” “conclude,” and “conclusion”

Ordinary talk of what has been “inferred” is normally talk of a new conclusion that is the result of inference, or perhaps an old conclusion whose continued acceptance is appropriately “reinforced” by one’s reasoning. We do not normally describe the discarding of a belief as something inferred, unless the belief is discarded as a result of accepting its negation or denial. But there are cases in which reasoning results in ceasing to believe something previously believed without believing its negation. In such a case it is somewhat awkward to describe the result of internal reasoning in terms of what has been inferred. Similarly, when reasoning leads one to discard something one previously accepted, it may be awkward to talk of the “conclusion” of the reasoning.

It might be said (misleadingly I think) that the “conclusion” of one’s reasoning in this case is to stop believing (or intending) X or, maybe, that one is (or ought) to stop believing or intending X . And, although it is syntactically awkward to say that what is “inferred” in such a case is to stop believing (or intending) X (because it is syntactically awkward to say that Jack inferred to stop believing (or intending) X), it might be said (again misleadingly I think) that what is “inferred” is that one is (or ought) to stop believing (or intending) X .

These (in my view, deplorable) ways of talking might also be extended to internal reasoning that leads to the acceptance of new beliefs or intentions. It might be said that the “conclusion” of one’s reasoning in such a case is to believe (or decide to) Y or that one is (or ought) to believe (or decide to) Y and that what one “infers” is that one ought to believe (or decide to) Y .

One of these ways of talking might seem to imply that all internal reasoning is practical reasoning, reasoning about what to do, to stop believing (or intending) X or to believe (or decide to) Y . The other way of talking might seem to imply that all reasoning is theoretical reasoning, reasoning about what is the case, it is the case that one is (or ought) to stop believing (or intending) X or it is the case that one ought to believe (or decide to) Y . Neither reduction is plausible given certain differences between theoretical and practical reasoning.

6 Characteristics of Theoretical and Practical Reasoning

Internal practical reasoning is concerned with what to do and internal theoretical reasoning is concerned with what is the case. Internal practical reasoning is reasoning that in the first instance is apt to modify one’s decisions, plans, or intentions; internal theoretical reasoning is reasoning that in the first instance is apt to modify one’s beliefs (“apt” because of limiting cases in which reasoning leaves matters as they are without any effect on one’s beliefs or intentions).

Of course, much internal reasoning is a mixture of practical and theoretical reasoning. One reasons about what is the case in order to decide what to do and one’s decision to do something influences what one believes will happen.

6.1 What internal practical and theoretical reasoning have in common

Internal reasoning of both sorts can be goal directed, conservative, and coherence seeking. It can be directed toward responding to particular questions; it can seek to make minimal changes in one's beliefs and decisions; and it can try to avoid inconsistency and other incoherence and attempt to make one's beliefs and intentions fit together better.

So, for example, one might seek to increase the positive coherence of one's moral views by finding acceptable moral principles that fit with one's opinions about particular cases and one might try to avoid accepting moral views that are in conflict with each other, given one's nonmoral opinions. In other words, one might attempt to reach a more or less wide "reflective equilibrium" (Rawls 1971). I will shortly discuss empirical research specifically directed to this model of internal reasoning.

6.2 How internal practical and theoretical reasoning differ

There are a number of ways in which internal practical and theoretical reasoning differ, having to do with wishful thinking, arbitrary choices, and direction of fit.

First, the fact that one wants something to occur can provide a reason to decide to make it occur but not a reason to believe it has occurred. Wishful thinking is to be pursued in internal practical reasoning but avoided in internal theoretical reasoning.

Second, internal practical reasoning can and often must make arbitrary choices, where internal theoretical reasoning should not. Suppose there are several equally good routes to where Mary would like to go. It may well be rational for her arbitrarily to choose one and follow it and it may be irrational for her not to do so. On the other hand Bob might be justified in believing that she is either taking route *A* or route *B* without being justified in believing that she is taking route *A* and without being justified in believing she is taking route *B*.

A third difference is somewhat difficult to express, but it has to do with something like the "direction of fit." Internal theoretical reasoning is part of an attempt to fit what one accepts to how things are. Internal practical reasoning is an attempt to accept something that may affect how things are. Roughly speaking, theoretical reasoning is reasoning about how the world already is and practical reasoning is reasoning about how if at all to change the world. Evidence that something is going to happen provides a theoretical reason to believe it will happen, not a practical reason to make it happen (Hampshire 1959). This way of putting things is inexact, however, because changes in one's beliefs can lead to changes in one's plans. If Mary is intending to meet Bob at his house and then discovers that Bob is not going to be there, she should change her plans (Harman 1976).

7 Against Reduction

I earlier mentioned proposals to reduce one type of internal reasoning—theoretical or practical—to the other, treating one type as a special case of the other. Some proposals (Levi 1967) take

what I am calling internal theoretical reasoning to be a special case of what I am calling internal practical reasoning, namely, the special case of reasoning to a decision about what to believe or stop believing. Others (Nagel 1970) take what I am calling internal practical reasoning to be a special case of what I am calling internal theoretical reasoning, namely, reasoning to a belief about what one *ought* to (or *may*) do.

These reductionist proposals have difficulty accounting for the differences between internal theoretical and practical reasoning. Levi's reduction of internal theoretical reasoning to practical reasoning would seem to entail that there is nothing wrong with arbitrary choice among equally good theoretical conclusions. Nagel's reduction of internal practical reasoning to internal theoretical reasoning simply denies that there is such a thing as internal practical reasoning in the sense of reasoning that results directly in decisions to do things and otherwise potentially modifies one's plans and intentions. Since neither reduction seems plausible to me, I continue to suppose that internal theoretical and practical reasoning are different if related kinds of reasoning.

8 Conscious and Unconscious Reasoning

Although some accounts of moral judgment identify reasoning with conscious reasoning (Haidt 2001), most psychological studies of reasoning have been concerned with unconscious aspects of reasoning. For example, there has considerable controversy about the extent to which reasoning about deduction makes use of deductive rules (Braine and O'Brien, 1998; Rips, 1994) as compared with mental models (Johnson-Laird and Byrne, 1991; Polk and Newell, 1995). All parties to this controversy routinely suppose that such reasoning is not completely conscious and that clever experiments are required in order to decide among these competing theories.

Similarly, recent studies (Holyoak and Simon, 1999; Simon 2001, 2004; Thagard 1989, 2000) investigate ways in which scientists or jurors reason in coming to accept theories or verdicts. These studies assume that the relevant process of reasoning (in contrast with its products) is not available to consciousness, so that evidence for theories of reasoning is necessarily indirect.

Actually, it is quite unclear that one is ever conscious of the activity of internal reasoning rather than of some of its intermediate and final upshots. Lashley (1958) famously asserted that "No activity of mind is ever conscious."

To be sure, people are conscious of (aspects of) the external discussions or arguments in which they participate and they can consciously imagine participating in such discussions. But that is not to say that they are conscious of the internal processes that lead them to say what they say in those discussions.

Since external arguments are expressed in words, imagined arguments will be imagined as expressed in words. This does not imply that internal reasoning is itself ever in words as opposed to being reasoning about something expressed in words (Ryle 1979) and does not imply that internal reasoning is ever conscious. When theorists refer to "conscious reasoning" they may be referring either to such externally expressed or imagined arguments or to other upshots of internal reasoning.

Gibbard (1990) and Scanlon (1998) suggest that moral thinking is at least sometimes concerned with finding ways of acting that can or could be justified to others. Internal moral reasoning might then involve thinking about external moral reasoning to others and so might always or typically involve conscious envisioning of such external reasoning. But internal reasoning about such external reasoning need not itself be conscious.

9 Parallel Constraint Satisfaction and Reflective Equilibrium

I mentioned earlier the idea that internal reasoning might take the form of making mutual adjustments to one's beliefs and plans, in the light of one's goals, in pursuit of what Rawls calls a "reflective equilibrium." Thagard (1989, 2000) develops computational models of this process using parallel "constraint satisfaction."

Thagard (1989) uses this idea to model the reasoning of jurors trying to assess the guilt of someone in a trial. The model makes certain predictions. For example, a juror might begin with a view about the reliability of a certain sort of eye-witness identification, a view about whether posting a message on a computer bulletin board is more like writing something in a newspaper or more like saying something in a telephone conversation, and so forth. Suppose the case being decided depends in part on an assessment of such matters. Then Thagard's model predicts that a juror's general confidence in this type of eye-witness identification should increase if the juror judges that in this case the testimony was correct and should decrease if the juror judges that in this case the testimony was not correct. The model predicts a similar effect on the juror's judgment about what posting on a computer network is more similar to, and so forth. The model also predicts that, because of these effects, the juror's resulting reflective equilibrium will lead to the juror's being quite confident in the verdict he or she reaches.

Experiments involving simulated trials confirm this prediction of Thagard's model (Simon 2004). In these experiments, subjects are first asked their opinions about certain principles of evidence about certain sorts of eyewitness identifications, resemblances, etc. Then they are given material about difficult cases involving such considerations to think about. The subjects' final verdicts and their confidence in their verdicts and in the various principles of evidence are recorded.

One result is that, as predicted, although subjects may divide in their judgment of guilt at the end, with some saying the defendant is guilty and others denying this, subjects are very confident in their judgments and in the considerations that support them. Furthermore, also as predicted, there are also changes in subjects' judgments about the value of that sort of eye-witness identification, about whether posting on a computer bulletin board is more like writing in a newspaper or having a private conversation, and so forth.

The model implies that judgments in hard cases are sometimes fragile and unreliable under certain conditions. When there is conflicting evidence, there is considerable tension among relevant considerations, just as there is a certain sort of tension among the nodes representing vertices in the Necker cube problem. If some nodes acquire even slightly increased or decreased excitation, the relevant inhibitory and excitatory connections can lead to changes in the excitation of other nodes in a kind of chain reaction or snowballing of considerations leading to a clear verdict, one way or

the other, depending on the initial slight push, just as happens in one's perception of a Necker cube.

After the Gestalt shift has occurred, however, the case may seem quite clear to the juror because of ways the juror's confidence has shifted in response to the positive and negative connections between nodes.

One upshot of this is that the slight errors in a trial that look like "harmless errors" can have a profound effect that cannot be corrected later by telling jurors to ignore something. By then the ignored evidence may have affected the excitation of various other items in such a way that the damage cannot be undone. Similarly, the fact that the prosecution goes first may make a difference by affecting how later material is evaluated.

This fragility of reflective equilibrium casts doubt about using the method of reflective equilibrium to arrive at reliable opinions.

This sort of problem has been noted in discussions of Rawls' (1971) claim that justification of views about justice consists in getting one's judgments into reflective equilibrium. It is sometimes suggested that the problem might be met by trying to find a "wide" rather than a "narrow" reflective equilibrium, where that involves not only seeing how one's current views fit together but also considering various other views and the arguments that might be given for them and trying to try to avoid the sorts of effects that arise from the order in which one gets evidence or thinks about an issue (Daniels, 1979). One needs to consider how things would have appeared to one if one had gotten evidence and thought about issues in a different order, for example. In this way one tries to find a robust reflective equilibrium that is not sensitive to small changes in one's starting point or the order in which one considers various issues.

Experimenters have shown that if subjects acting as jurors are instructed to try for this sort of wide robust reflective equilibrium, they are less subject to the sorts of effects that occur when they are not so instructed (Simon, 2004). But the effects do not go away completely.

10 The Role of Moral Principles in Moral Reasoning

Sometimes moral conclusions about particular cases appear to be based on general moral principles. But, according to moral particularism, particular moral judgments need not be based on general principles of any sort (Dancy 1993; for critical discussion see e.g. Sinnott-Armstrong 1999, Hooker and Little 2000, Kihlbom 2002, Väyrynen, 2004, McKeever & Ridge, forthcoming).

Some moral particularists suppose that some moral judgments are (or are like) direct perceptions of (apparent) moral truth (McDowell 1979). That may appear to conflate moral judgments with aesthetic judgments. Aesthetic judgments do often rest on perception and seem not to depend on general principles (Sibley 1959, Isenberg 1949). Hampshire (1954) argues that this represents a crucial difference between aesthetic and moral judgments.

The issue is delicate, because a tendency to have intuitions or emotional reactions might involve an

implicit acceptance of a corresponding general moral principle that one may not be able formulate in words (Hare 1952, pp. 56-78). The principle might be implicit and practical rather than explicit and theoretical, representing know-how rather than knowledge that. In this view, one has such principles whether or not one can formulate them as long as one acts or reacts from the relevant dispositions.

Can such “implicit principles” be modified by reasoning in the way that explicit principles can be? It may seem that explicit principles are more easily changed whereas implicit principles are like habits and so hard to change. But it is not clear that there really is such a difference. On the other hand, at least according to Hare, the relevant behavioral dispositions sometimes change quickly and permanently on the basis of experience. Furthermore, some explicit opinions are hard to get rid of in the face of evidence against them (Lord, Ross, and Lepper, 1979).

Discussions of moral particularism often appear to assume that a person’s general moral principles must be like the familiar principles that people actually state and use in everyday argument. This conflicts with Hare’s (1952) view and with those (e.g., Rawls 1971; Mikhail 1996, forthcoming; Dwyer 1999; Harman 1999; Hauser 2006) who have proposed an analogy with linguistics. Recent progress in linguistics depends on supposing that relevant principles not be restricted to those with which ordinary speakers are familiar.

Such a restriction for moral principles may make sense if the issue is limited to the sorts of principles invoked in external reasoning with others rather than the principles that might be involved in internal reasoning.

11 Hard Cases

Sometimes when a moral question arises, it’s easy to answer because it clearly falls under a principle one accepts and there is nothing special about the case. But other cases are not like that. How do people reason about these other hard cases?

Sometimes, people try to decide in a way that minimizes implications for certain other hard cases.

This happens for external reasoning in certain legal contexts in which a court has to reach a decision in a case not covered by previous legal principles. The court will sometimes try to decide the issue narrowly so as to minimize the impact on decisions about other hard cases. Members of the US Supreme Court occasionally disagree about whether the majority decision in a given case has implications for various other possible cases, with those in the majority arguing that it has no such implications. On such occasions it is sometimes taken to be a virtue in a decision that the decision leaves other possible nearby cases maximally undecided.

For example, in a 2003 case, *Lawrence et al. v Texas*, the U. S. Supreme Court majority ruled that an anti-sodomy law unconstitutionally discriminates against homosexuals. In a concurring opinion, Justice Sandra Day O’Connor argued that the rationale of the Court’s decision would not also apply to laws prohibiting same sex marriages. In dissent, Justice Antonin Scalia disagreed. Both opinions took for granted that it is undesirable to decide the issue then under consideration

in a way that would also decide that other issue.

The principle behind such legal reasoning resembles one version of what in a very different context Vapnik (1998, 2000) calls “transduction.” Transductive reasoning is based on the information that certain cases are hard cases and that one or more particular hard cases have come up to be decided. One version of transduction seeks to extend previous decisions to the new cases in a way that minimizes decisions about “hard cases.” Learning machines using transduction outperform other systems (Joachims 1999, Weston et al. 2003). It is an interesting question whether human reasoning uses anything like transduction.

It is of course true that courts engage in external public reason. But similar considerations may apply to some internal reasoning. For one thing, the way a person classifies a new case often does depend on what other new cases are to be classified. Psychologists often suggest this illustrates “irrationality” in human thinking, a “framing effect,” perhaps. But it may indicate instead that people sometimes reason transductively.

Redelmeier and Shafir (1995) discuss the following example. Suppose that a certain painful condition can be alleviated by either of two medicines, each of which has different side effects. If only one is available, doctors tend to prescribe it for this condition. If both are available, doctors tend to prescribe neither, presumably because they have difficulty deciding between them. Similarly, customers who come upon a display of six jams, are more apt to decide to purchase one of the jams than are customers who come upon a display of twenty four jams including those six (Iyengar and Lepper 2000). The task of deciding between all twenty four jams is presumably too difficult to be worth the cost.

Not being willing to take the trouble to choose among many jams (and similar examples discussed by Schwartz, 2004) seems quite reasonable, much more reasonable than the doctors’ not taking the trouble to choose between two rather different pain killers, thereby leaving the patients to suffer. It is unclear whether either of these cases or any other cases are appropriately modeled as instances of transduction.

12 The Extent of Internal Reasoning

I have distinguished external reasoning, involving discussion with other people, from internal reasoning processes. I discussed how internal reasoning is apt to lead to changes in one’s beliefs, intentions, desires, and possibly other psychological attitudes, even emotions. This leaves the question how to distinguish those internal processes that are instances of reasoning from those that are not. This is partly a purely verbal issue but also partly a substantive theoretical issue about the most fruitful way of envisioning internal reasoning.

Internal reasoning is a process that typically (if not always) involves a change of some sort. But what changes are part of the reasoning process and what are merely partial causes or results of changes that are part of the reasoning process? Practical reasoning typically involves changes in what one plans or intends to do and can lead one to act in certain ways—for example, to walk somewhere. It is plausible to us that changes in plans and intentions can be part of practical

reasoning, whereas walking and other actions are merely intended results of the reasoning.

What about desires and emotions? These are plausibly counted as part of internal reasoning to the extent that there are reasons or grounds for having them and they provide reasons or grounds for other attitudes. It makes sense to suppose that desires play a role in practical reasoning to the extent that desires provide reasons or grounds for taking one course of action rather than another and to the extent that there can be reasons or grounds for desiring one thing rather than another. Desires have a rational effect on conclusions of practical reasoning (including decisions about what theoretical issues are worth pursuing). Desires are rationally affected by reasoning which indicates reasons to want one thing rather than. One can have grounds to desire something as a means to something else one desires, for example. And one can have grounds to stop desiring something or concluding that it will not after all result in what one thought it would result in.

Are there similar reasons to think that reasoning can involve emotions? I am not sure about this, my way of thinking about internal reasoning makes this a natural question to ask. Are the processes by which emotions arise even processes of reasoning? Just as desires can play a role in reasoning, perhaps other feelings can also play a role—not a distorting role but a rational role in something like the way in which desires play such a role.

Reasoning can change what one accepts, including what one desires, and it can change one's emotions. Anger can depend on beliefs that provide grounds for the anger and, if the beliefs are given up, the grounds for anger are gone and the anger modified. An emotion might serve as part of one's grounds for accepting a moral judgment: "I am feeling disgust, so something must be wrong." In this instance, it is perhaps a belief about an emotion that functions in reasoning, but emotions may also figure in moral thinking in something like the way in which desires figure in practical reasoning, so that emotions themselves, and not just beliefs about them, would be inputs to moral decisions. I do not know whether there is a good way to think about the possible relationship between emotions and moral reasoning, but I do think it is something worth taking seriously.

13 Concluding Summary

I began by noting that my topic is the process, not the product, of reasoning. I observed that formal logic, probability and decision theory are not by themselves theories of reasoning in this sense. I distinguished internal and external reasoning and noted the difficulty of distinguishing reasoning from certain other internal computational processes. I pointed to differences between theoretical and practical reasoning that tell against trying reduce one to the other. I agreed with psychologists who find that the process reasoning tends to be unconscious, although its products may be conscious. I described Thagard's parallel constraint satisfaction simulation of reasoning toward a reflective equilibrium and some psychological evidence in support of his model. I discussed the role of principle in internal reasoning and the debate over moral particularism, a debate that makes the odd assumption that the relevant principles must be the sorts of familiar principles people appeal to in discussion. I discussed one way in which people deal with certain "hard cases." I ended with some speculation about the role of emotion in reasoning.

14 References

- Braine, M. D. S. and O'Brien, D. P. (1998). *Mental Logic*. Mahwah, NJ: Lawrence Erlbaum Associations.
- Daniels, N. (1979). "Wide Reflective Equilibrium and Theory Acceptance in Ethics." *Journal of Philosophy* 76: 256-82.
- Dancy, J. (1993). *Moral Reasons*. Oxford: Blackwell.
- Dwyer, S. (1999). "Moral Competence." *Philosophy and Linguistics*, edited by Kumiko Murasugi and Robert Stainton. Boulder, CO: Westview Press, 169-190.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings : A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Haidt, J. (2001). "The emotional dog and its rational tail: A social intuitionist approach to moral judgment." *Psychological Review* 108, 814-834.
- Hampshire, S. (1954). "Logic and Appreciation." In Elton, W. (ed.), *Aesthetics and Language*. Oxford: Blackwell, 161-169.
- Hampshire, S. (1959). *Thought and Action*. London: Chatto and Windus.
- Hare, R. M. (1952). *Language of Morals*. Oxford: Oxford University Press.
- Harman, G. (1976). "Practical Reasoning," *Review of Metaphysics* 29, 431-463.
- Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
- Harman, G. (1999). "Moral Philosophy and Linguistics" in Klaus Brinkmann, ed., *Proceedings of the 20th World Congress of Philosophy, Volume 1, Ethics*. Bowling Green, OH: Philosophy Documentation Center, 107-15.
- Harman G. and Kulkarni, S. (2007). *Reliable Reasoning: Induction and Statistical Learning Theory*. Cambridge, MA: MIT Press.
- Hauser, M. (2006). *Moral Minds: How Nature Designed Our Universal Sense of Right and Wrong*. New York: HarperCollins.
- Holyoak, K. J. and Simon, D. (1999). "Bidirectional Reasoning in Decision Making by Constraint Satisfaction," *Journal of Experimental Psychology: General* 128: 3-31.
- Hooker, B. and Little, M. (2000). *Moral Particularism*. New York, Oxford University Press.
- Isenberg, A. (1949). "Critical Communication," *Philosophical Review* 58, pp. 330-344.

- Iyengar, S. S. and Lepper, M. R. (2000). "When Choice Is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology* 79: 995-1006.
- Joachims, T. (1999). "Transductive Inference for Text Classification Using Support Vector Machines." In I. Bratko and S. Dzeroski, editors, *Proceedings of the 16th International Conference on Machine Learning*, pp. 200-9. San Francisco: Morgan Kaufmann.
- Johnson-Laird, P. N., and Byrne, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kihlbom, U. (2002). *Ethical Particularism*. Stockholm Studies in Philosophy 23. Stockholm: Almqvist and Wiksell.
- Lashley, K. S. (1958). "Cerebral organization and behavior." In Solomon, H.C., Cobb, S., and Penfield, W. (Eds.) *The Brain and Human Behavior*, Vol. 36. Association for Research in Nervous and Mental Disorders, Research Publications. Baltimore: Williams and Wilkin. Pp. 1-18.
- Levi, I. (1967). *Gambling with Truth*. New York: Knopf.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology*, 37, 2098-2109.
- McDowell, J. (1979). "Virtue and Reason." *The Monist* 62: pp. 331-50.
- McKeever, S. & Ridge, M. (2006). *Principled Ethics: Generalism As a Regulative Ideal*. Oxford: Oxford University Press.
- Mikhail, J. (1996). "The Moral Faculty." Unpublished workshop presentation at MIT.
- Mikhail, J. (forthcoming). *Rawls Linguistic Analogy*. New York: Cambridge University Press.
- Nagel, T. (1970). *The Possibility of Altruism*. Oxford: Oxford University Press.
- Polk, T. A., and Newell, A. (1995). "Deduction as Verbal Reasoning," *Psychological Review*, 102, 533-566.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Redelmeier, D. A. and Shafir, E. (1995). "Medical Decision Making in Situations that Offer Multiple alternatives," *Journal of the American Medical Association* 273: 302-5.
- Rips, L. J. (1994). *The Psychology of Proof*. Cambridge, MA: MIT Press.
- Ryle, G. (1979). *On Thinking*. Totowa, NJ: Rowman and Littlefield.

- Scanlon, T. M., (1998). *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- Schwartz, B., (2004). *The Paradox of Choice: Why More Is Less*. New York: HarperCollins.
- Sibley, F. (1959). "Aesthetic Concepts," *Philosophical Review* 67, pp. 421-450.
- Simon, D., et al. (2001). "The Emergence of Coherence over the Course of Decision Making," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 27: 1250-1260.
- Simon, D., (2004). "A Third View of the Black Box: Cognitive Coherence in Legal Decision Making." *University of Chicago Law Review* 71: 511-586.
- Sinnott-Armstrong, W. (1999). "Varieties of Particularism," *Metaphilosophy* 30: 1-12.
- Thagard, P. (1989). "Explanatory Coherence," *Behavioral and Brain Science* (1989) 12: 435-467.
- Thagard, P. (2000). *Coherence in Thought and Action*. Cambridge, MA: MIT Press.
- Vapnik, V. (1998). *Statistical Learning Theory*. New York: Wiley.
- Vapnik, V. (2000). *The Nature of Statistical Learning Theory*, second edition. New York, Springer.
- Väyrynen, P. (2004). "Particularism and Default Reasons," *Ethical Theory and Moral Practice* 7: 53-79.
- Weston, J., Perez-Cruz, F., Bousquet, O., Chapelle, O., Elisseeff, A., and Scholkopf, B. (2003). "KDD Cup 2001 Data Analysis: Prediction of Molecular Bioactivity for Drug Design-Binding to Thrombin." *Bioinformatics*.