

PRACTICAL RATIONALITY IS A PROBLEM IN THE PHILOSOPHY OF MIND

Timothy Schroeder
Ohio State University

The philosophy of mind encompasses a familiar set of topics: consciousness, intentionality, mental causation, emotion, whatever topics in psychology happen to capture our interest (concepts, mindreading . . .), and so on. There is a topic deserving of addition to this list, a topic that should be receiving regular attention from philosophers of mind but is not: practical rationality. The philosophy of mind bears directly upon what can be called the ‘meta-theory’ of practical rationality, and meta-theories of rationality likewise impose constraints on the philosophy of mind. Philosophers of mind who do not engage with work on practical rationality risk developing theories that are blind to important theoretical demands. They also risk allowing ethicists—who *have* been at work on the meta-theory of practical rationality—to unilaterally set their theoretical agendas. It is therefore time to get to work.

1. Practical Rationality

It will help to begin with a distinction that belongs entirely to ethics: that between normative ethical theories and meta-ethical theories.

Normative ethical theories are theories of what is morally right and morally wrong.¹ Ideally, normative ethical theories identify what is right and wrong in non-moral, non-evaluative terms. For example, utilitarianism is a normative ethical theory that identifies what is right or wrong in terms of states of pleasure, while Kantianism is a normative ethical theory that identifies what is right or wrong in terms of the consistency of the will’s contents.²

Meta-ethical theories are theories of how the non-moral, non-evaluative things identified by normative ethical theories come to have their ethical

significance: how the “is” facts are related to the “ought” facts.³ If it is the case that a certain action maximizes expected pleasure, that is an “is” fact. If utilitarianism is correct, this “is” fact also adds up to⁴ an “ought” fact: the action is the one that ought to be performed. What makes it true that the “is” adds up to the “ought”? Perhaps it is the case that a divine being has ordered us to take actions that maximize expected pleasure. This is one possible meta-ethical explanation for why the “is” adds up to the “ought.” Or perhaps it is the case that there is a local social practice that favours acting so as to maximize expected pleasure. This is a different possible meta-ethical explanation. And, of course, divine command theory and cultural relativism are just two of the meta-ethical theories available. Equally obvious is that there can also be a divine command theory of Kantian normative ethics, or a culturally relativistic theory of Kantian normative ethics, and so on. Meta-ethical theorizing is largely independent of normative ethical theorizing.

The same distinction can be applied to the theory of practical rationality. One can focus on a normative theory—which actions are the rational ones?—or on a meta-theory—how do the “is” facts add up to the fact that some actions are *rational* and others *irrational*? A familiar answer to the normative theoretical question regarding practical rationality is provided by decision theory. Standard decision theory tells us that the rational action is the action that maximizes expected satisfaction of preferences. Other approaches to the normative theoretical question have been provided by ethicists. Kantians, for instance, characteristically hold that rational action is constrained by considerations of the universalizability of the maxim on which action proceeds: when the maxim is not universalizable, then it is irrational to act on it even if so acting would maximize expected satisfaction of preferences.⁵

Focus on the meta-theory of practical rationality reveals a less satisfying state of affairs. As it happens, the meta-theoretical question is rarely asked, and no answer that has been offered has gained particularly wide acceptance. Furthermore, most of the answers on offer have been developed by ethicists.⁶ As will be seen, there is good reason for philosophers of mind to join this field of inquiry.

2. Dismissive Meta-Theories of Practical Rationality

One reason for the dearth of meta-theories might be that there is no true and substantial meta-theory of practical rationality. Another reason might be that the meta-theory is trivial to provide. In this section, both possibilities are considered.

One answer to the meta-theoretical question is provided by the expressivism of Allan Gibbard in *Wise Choices, Apt Feelings*.⁷ Gibbard holds that to call an action rational is just to express approval of it. As a result, Gibbard’s

theory is one that ultimately denies that there are interesting “ought” facts about practical rationality: there is what we approve of, but approval is not held to be something that can take an “is” fact and turn it into an “ought” fact. Rather, approval licences certain forms of “ought” speech, nothing more.⁸

There is already a substantial critical literature on Gibbard, and I will simply note it; the literature is too large and too technical to engage properly here.⁹ If Gibbard is right, there is indeed no problem to be solved. But there is reason to think Gibbard is not right. At the very least, it certainly seems that claims about the rationality of actions are claims about the world, and not mere expressions of approval. If this appearance is correct, then this paper has a point. What follows starts from the assumption that this paper does indeed have a point, but this will simply have to be an assumption.

Another objection to the present project holds that there is no meta-theoretical question to be answered. Once one has a normative theory of practical rationality, one has all the theory one needs. Such a view might be found in work by Steve Stich, who has argued (if not in these terms) that the meta-theory of *theoretical* rationality should appeal to practical rationality (to the pragmatic interests one has in believing truths, and so on), without treating the meta-theory of *practical* rationality as equally problematic.¹⁰ Perhaps there is no comparable problem regarding practical rationality because there just is no meta-theory of practical rationality to be stated. Whatever Stich’s view, the objection is a reasonable one to consider. There is no need, one might say, to explain why a person *ought* to act rationally. Irrational action is its own punishment, as it were.

In response, it is important to focus on the basic facts. Rational and irrational actions, like moral and immoral actions, are events that are (at least) *described* as inherently evaluative in nature. This distinguishes them from other events. There is nothing inherently evaluative in describing a star as going supernova. This is not a type of event that is, of necessity, good or bad, right or wrong, as it should be or as it should not, functioning or malfunctioning. But rational actions *are* described as inherently evaluative in nature. They are said, in one respect at least, to be good, right, functioning properly, or as they ought to be: they are *rational*, as opposed to *irrational*. This way of talking and thinking about rationality needs explaining. And the most plausible explanation, setting expressivism aside, is that we talk and think this way because it is true. Of the various properties that rational actions possess, one is that they are positively evaluated by a certain standard: the standard of practical rationality itself. This fact is in need of explaining.

Of course, the evaluative nature of practical rationality “needs” explaining purely from a theoretical perspective, and so there is some truth in the thought that one does not need to explain to a person why she ought to act rationally. The explanation of why rational actions possess evaluative properties, and do so of necessity given that they are rational, is not an

explanation that has practical import. But the explanation of why certain neural states possess consciousness properties, and do so of necessity (if they do), is not an explanation that has practical import either. It is nonetheless of great theoretical interest. Likewise for the meta-theory of practical rationality. The theoretical question does not go away just because an answer to it is neither required to encourage rational behaviour in people nor useful for that purpose.

A different objection to the project of providing a meta-theory for practical rationality holds that the project is trivial, because irrational action is self-defeating and because the badness (wrongness...) of self-defeating action is trivial.

This objection is worthy of more serious consideration, because it recognizes that there is a phenomenon to be explained and proposes an explanation. Understanding what is wrong with the objection requires going deeper into the meta-theory of practical rationality.

Suppose one has a normative theory of practical rationality in hand: the non-evaluative properties that make an action rational are such-and-such. Then say that an action has non-evaluative properties PR, for short. It is thus supposed that having PR adds up to the property of being rational, and that being rational is an evaluative property. What is needed is an explanation of why this is so. That is, what is needed is an explanation of why the “is” of having PR adds up to the “ought” of being rational. And the proposed explanation is that lacking PR would make the agent’s action self-defeating. This much is now clear. But how is “lacking PR would make the agent’s action self-defeating” a good explanation?

The demands on a good explanation here are stringent, as they are anytime one is trying to explain how having properties A, B, and C adds up to having property X. Consider work on consciousness. Philosophers proposed that having property A (the property of containing firing C-fibres) added up to having property X (the property of being in pain). This simple identity theory is now recognized as fundamentally deficient on account (at least) of its not providing an adequate explanation of why having property A adds up to having property X. Or, to put it in terms more familiar from the literature on consciousness, identity theory does not show us that we cannot coherently conceive of a being that has the property of containing firing C-fibres while lacking the property of being in pain.¹¹ If there were an outstanding explanation of why having the first property added up to having the second, it would not be possible to coherently conceive of a being having one property but not the other: it would be explained why the having the one property adds up to having the other. The explanation would tell us how to think of C-fibres and of pain such that their separation would be rendered inconceivable. But, unfortunately for identity theory, there is no such explanation. And unfortunately for materialism in the philosophy of mind, there is a certain scepticism in the literature regarding the possibility

of explaining how any natural, non-mental set of properties A, B, C . . . could add up to the property of seeing X or tasting Y.¹²

The problem for the meta-theory of practical rationality is (at least) to show how it is inconceivable that one could have PR but fail to have oneself or one's action evaluated as good, right, doing what one ought, or something similar, in the relevant sense. The nature of PR and the nature of rationality of action must be shown to be such that one adds up to the other. And talk of the self-defeatingness of actions that lack PR does not suffice for this purpose. Self-defeatingness is a property an action possesses in virtue of the action's success at bringing about its end, A, entailing¹³ that although one desires P, Q, R . . . more than one desires A, one will not (now) be able to bring it about that P, Q, R . . .¹⁴ For instance, the respect in which playing computer games all day is self-defeating for me is that I desire to do the laundry, and if I play computer games all day then that entails I will not do the laundry, but I want more to do laundry than to spend all day playing computer games. Given all of this, it certainly seems *true* that self-defeating actions are irrational. But since self-defeat seems to be well characterized in the non-evaluative terms just used, something vital has not been done. The idea of self-defeat has not been shown to add up to something evaluative. The philosopher (such as Gibbard) who claims to be able to conceive of self-defeat without conceiving of an "ought" coming to exist out of the facts of what "is" has not been shown to be making a conceptual error. The central meta-theoretical task has not, therefore, been accomplished.

A related dismissive strategy fails for similar reasons. Suppose one holds that there is a kind of inconsistency instantiated when one fails to instantiate PR. One might then hold that this inconsistency is a negative evaluative state, and so explains how failing to instantiate PR adds up to acting in a way one ought not to act. The problem is that inconsistency is not inherently evaluative, any more than self-defeat is. Many propositions are inconsistent with one another, but this is not (in of itself) something good or bad, right or wrong. Likewise, many propositions held as contents of attitudes are inconsistent without this inconsistency being in any way evaluatively significant: your beliefs might not be (theoretically) consistent with mine, and your intentions might not be (practically) consistent with my attitudes, but none of this matters one bit, evaluatively speaking. There is nothing wrong with your beliefs for not being (theoretically) consistent with mine, and there is nothing wrong with your intentions for failing to be (practically) consistent with my attitudes. There is, of course, something wrong with my intentions when they are practically inconsistent with *my* attitudes. But mere consistency is not the heart of the explanation of what is wrong, for mere consistency itself is not something that generates evaluative facts out of non-evaluative ones.¹⁵

Although this section has only reviewed a few possible dismissive approaches to practical rationality, perhaps the reader is ready to agree that

there is a real problem to be solved by the meta-theory of rationality, and is ready to turn to consideration of non-dismissive attempts to solve it.

3. Substantive Meta-Theories of Practical Rationality

Philosophers of mind have several tools in their toolboxes for getting from “is” facts to “ought” facts in a way that makes the existence of the “is” without the “ought” genuinely inconceivable: in a way that genuinely explains how the “is” facts add up to the “ought” facts. In this section, I consider whether use of these tools will help to get from the “is” facts of PR (the correct normative theory of practical rationality) to the “ought” facts of practical rationality and irrationality.

Begin with social practices. Neo-Wittgensteinians such as Wilfred Sellars and Robert Brandom have argued for a role for social practices in theorizing the mind.¹⁶ And both have given a role to social practices precisely because both argue that social practices can generate important “ought” facts.¹⁷

Whatever one thinks of the specifics of Sellars’ or Brandom’s theorizing in the philosophy of mind, it is probably correct that certain social practices take “is” facts and turn them into “ought” facts: the tool is sound, whatever one makes of the particular uses to which it is put. Consider that restaurants generally set forks on the left of plates and knives on the right of plates. This is also where the cutlery ought to be: a restaurant that bucked convention (say, out of cultural ignorance) would be putting the cutlery in the wrong place. And it seems plausible that it is some fact about social practices that explains how this “is” becomes an “ought.” Certainly it appears that if the social practices were different that would change the “ought” facts, which suggests that social practice is quite powerful in this regard. But one can go deeper and see this as conceptually necessary. Certain social practices regulate particular domains, and when a practice regulates a domain, it generates a regulation. The regulation just is the standard against which the regulatory practice measures the domain in question, in determining what course of action to take in regulating the domain. In this way, social practices regulating domains make for regulations where none might previously have existed, just by the nature of regulatory practice. And regulations are evaluative standards.¹⁸

That the “is” facts of regulation generate “ought” facts regarding the domain regulated can be seen in the social world both on the broad scale of social regulation (exemplified by the little rules of etiquette) and on the smaller scale of individuals regulating domains on their own. To see this latter in action, imagine that I regulate Nicholas the cat’s scratchings: I do nothing to stop him from scratching most objects, but take various steps to stop him from scratching the armchairs. Then the regulation is that Nicholas may not scratch the armchairs: an evaluative standard is created. And this fits ordinary intuition about the case, which holds that Nicholas is not supposed

to scratch the armchairs under these conditions. Even if my regulation of his scratching is not especially efficacious, notice, it seems that the rule still exists—which is a further sign that there is indeed an “ought” at work here, and not a mere statistical expectation or the like.

In the light of the aforementioned observations, it seems reasonable to grant that facts about social practices of various sorts and scales can indeed serve as explanations of “ought” facts. But can social practices explain the “ought” facts of present interest, the “ought” facts about practical rationality? Two things seem required. First, it would have to be true that there is a social convention requiring people to act so as to instantiate PR (to maximize expected desire satisfaction, or . . .). And second, it would have to be true that the properties of the “ought” generated by such a social convention correspond to the properties of the “ought” of practical rationality.

It is perhaps true that there are social practices rewarding people who act so as to instantiate PR and punishing those who do not. People are prone to exasperation at others who act in self-defeating ways, even when little is at stake for the person who is exasperated. We generally encourage others to “get their acts together” and discourage others from senseless behaviour. Perhaps this adds up to an appropriate social practice regulating one another’s actions in a manner aimed at getting them to instantiate PR. This is not particularly powerful evidence, but suppose that there are such social practices. Then these social practices would make it true, as a matter of conceptual necessity, that there is a regulation, an evaluative standard, favouring instantiating PR in action. And this is just what the meta-theory requires. Is it as easy as this?

It is not so easy. The rightness (goodness, reasonableness . . .) of instantiating PR does not seem to depend on the existence of social convention. Even if I were the only person in the universe, alone on an island, my actions would be evaluated as right or wrong based on their instantiation of PR or failure to instantiate it. For instance, I might know that there are certain berries that, while delicious, will cause pain and fever a few hours after being eaten. And one day, while hungry, I might be tempted to eat just a few of these berries. But to do so would be irrational, and refraining rational. One’s actions can have these properties even in the absence of social convention. At least, so it seems to many, who would hold that we would have minds under these conditions. But grant, for the sake of the neo-Wittgensteinian, the idea that being enmeshed in the right social practices is required to have a mind at all. It still does not follow that the properties of the “ought” stemming from social practice are the right properties for the “ought” of practical rationality. For suppose it were true that social practices mandated, not PR, but something else (say, spontaneity) in action. Even so, spontaneous action that violated PR would not thereby be made into rational action: social practices are not enough to make spontaneity always reasonable, after all.

What goes for the social level also goes for the individual level: it is not necessary that an actor act to regulate her own behaviour in favour of instantiating PR in order for her to be rational. She might, for instance, hold a false normative theory of practical rationality, and attempt to get herself to instantiate some other pattern, rather than PR, in her actions. Doing so would not change the fact that she was trying to make herself act irrationally, foolishly, wrongly. This is the situation that the naive devotee of Ayn Rand can find herself in, for instance: regulating her own actions for selfishness in spite of the fact that acting selfishly is often not the rational course of action for her.¹⁹ The self-regulator, like society as a whole, can certainly create a regulation, and so an evaluative standard, and thereby create new evaluative facts. But these facts are powerless to change the facts about rational action, and so it is clear that rationality itself does not stem from social practice.

Thus, however useful this particular tool is for other purposes (explaining the existence of norms for utterances?), it is not useful for the meta-theory of practical rationality.

Turn now to natural selection. Ruth Millikan, David Papineau, and other teleosemanticists have argued for a role for natural selection in theorizing the mind. And they have given a role to natural selection precisely because it seems to play a role in turning “is” facts into “function” and “malfunction” facts—into specific “ought” facts, in other words.

Whatever one thinks of the specifics of teleosemantic theorizing in the philosophy of mind, it is probably correct that natural selection takes “is” facts and turns them into “ought” facts: the tool is once again sound, whatever one makes of the particular uses to which it is put. Consider that hearts pump blood. They also have the function of pumping blood: this is what a well functioning heart does. And it seems plausible, after about forty years of intensive work in the philosophy of biology, that hearts have their functions explained by their histories of natural selection.²⁰ And again, as with social convention, one can go deeper and see the power of natural selection to create its function “oughts” as conceptually necessary. Natural selection is a (non-intentional) design process, as Dennett in particular has emphasized,²¹ and things created to a design have the function of doing what they are designed to do.²² In this way, natural selection makes for evaluative standards where none might have previously existed, just by the nature of processes of design.

In the light of the aforementioned observations, it seems reasonable to grant that natural selection can indeed serve as an explanation of at least function-derived “ought” facts. But can natural selection explain the “ought” facts about practical rationality? As with social convention, two things seem required. First, natural selection would have to give people the function of producing actions that instantiate PR. And second, it would have to be true that the properties of the “ought” generated by natural selection correspond to the properties of the “ought” of practical rationality.

It is perhaps true that natural selection has given us, or structures within our brains, the function of producing actions that instantiate PR. At least, we seem reasonably good at instantiating PR when our brains are functioning properly, and things we think of as interfering with biological function (neural injury, drugs, lack of blood sugar) are familiar as causes of actions that do not instantiate PR. Suppose then that we do have the function of acting so as to instantiate PR. Then natural selection would make it true, again as a matter of conceptual necessity, that there is an evaluative standard favouring instantiation of PR. And this is just what the meta-theory requires. Is it as easy as *this*?

Again, it is not. The rightness (goodness, reasonableness . . .) of producing actions instantiating PR does not appear to depend on the existence of a history of natural selection any more than it depends on the existence of social practices. Even if we were, like Donald Davidson,²³ the children of an earthly swamp and a heavenly lightning bolt, we would be capable of acting rationally or irrationally: as we ought or as we ought not. At least, so it seems to many, who would hold that we would have minds under these conditions. But grant, for the sake of the teleosemanticist, the idea that a history of natural selection is required to have a mind at all. It still does not follow that the properties of the “ought” stemming from natural selection are the right properties for the “ought” of practical rationality. For suppose it were true that we had been naturally selected to sometimes have sex with biologically fit and fertile partners regardless of the social, emotional, and moral consequences of such sex. (This does not seem like an enormous supposition, truth be told.) Were this the case, that would not suffice to make such sex practically rational. PR would still be the correct normative theory of rationality, under such evolutionary conditions: biological imperatives are not enough to make a course of action truly reasonable.

Thus, however useful this particular tool is for other purposes (explaining the existence of mental representations?), it is no more useful for the meta-theory of practical rationality than social practices.

4. The “Ought” of Practical Rationality

So far this work has relied upon simple intuitions regarding the “ought” of practical rationality: it is not affected by foolish individual policies, it can survive the absence of a social network, and it is indifferent to the details of the natural selection of systems within our brains. It would be preferable to have something more principled to say about the evaluative standard in question. What are its general properties, and how can these guide a search for the correct meta-theory?

Three features of practical rationality seem most important. The first is that the “ought” of practical rationality applies to all the actions of all

the creatures capable of acting. The second is that the “ought” of practical rationality applies in virtue of the nature of action itself. And the third is that the “ought” of practical rationality requires conformity to PR.

The claim that the “ought” of practical rationality applies to all actions amounts to the claim that every action is one that is rational or irrational to some extent.²⁴ There is both intuitive and principled support for this idea. Intuitively, the domain of bodily movements we evaluate for rationality is the same as the domain of bodily movements that we consider to be actions, to be things done by the person as she moves her body. It is rational of Katie to drink coffee, given her circumstances, and drinking coffee is something she does: it is one of her actions. It is neither rational nor irrational of Katie to shiver in the cold, given her circumstances, but also her shivering just happens to her, and is not something she does: it is not one of her actions. And so on, for various cases. In a more principled vein, it can be said that to act more or less rationally is to act for better or worse practical reasons, and there are principled theories of action requiring that all actions be performed for some (better or worse) practical reason.²⁵ The upshot of both intuition and principle is that every action is in the domain of evaluation of practical rationality, and so the facts that make PR an “ought” must do so for every action.

The claim that the “ought” of practical rationality applies in virtue of the nature of action itself is more contentious. In ethics, for instance, there is no meta-ethical consensus that the “ought” of morality applies in virtue of the nature of moral action itself. That it does so is a famous Kantian contention,²⁶ but the contention is contradicted by both divine command theory (the “ought” exists in virtue of a specific divine command) and cultural relativism (the “ought” exists in virtue of an ongoing social practice), to name two. And if ethics is so divided, why should the meta-theory of practical rationality be any different? In defence of the claim, consider the failures of the two approaches considered in the previous section of this paper. The main problem with social practice was not that there might be action outside of society, but that it seems that social—even, individual—decision to regulate action is not the right source for the “ought” proper to action. Both are external to the nature of action itself. Society is, of course, external to the capacity of an individual to act. But so too is a policy set by an individual for herself, because such a policy is an *expression* of the individual’s capacity to act, rather than a component of what makes up her action capacities. If society sets a standard of behaviour for an individual, the individual can always ask, “but why should I care?” This is also true of a standard that an individual sets for her own behaviour. Having set it, the individual can ask, “but why should I care about doing things that way? Perhaps I am telling myself to act foolishly.” It is only a standard that is inherent to the nature of action itself—a standard that cannot be dismissed except by someone dismissing action itself—that can avoid this

sort of challenge. Much the same lesson comes from thinking about natural selection: that we are naturally selected to act in a certain way seems a fact of little moment, a fact that can be dismissed without action itself being dismissed. Only an “ought” maker inherent in action cannot be dodged by an actor. Though not a decisive argument by any means, there is some force to this.

Finally, the claim that the “ought” of practical rationality must require conformity to PR is obvious, but worthy of just a little attention on its own. In the search for things that can turn “is” facts into “ought” facts, it might seem promising to latch onto anything that can apparently do the hard work of getting from “is” to “ought.” But sometimes, however plausible it might be that the gap from “is” to “ought” has been bridged, it will be independently implausible that the “ought” on the far side of the bridge requires specifically PR. For example, consider the fact that the practically rational action is (according to many normative theories) the best action, with an approximate continuum of worse actions available. If this is right, then whatever explains the “ought” of practical rationality will have to be something that itself admits of degrees, or degrees of application, or degrees of satisfaction—something that will make sense of the notion of doing a better or worse job of satisfying it, if it is not perfectly satisfied. Not every way of generating an evaluative standard has this feature. A social practice in which all crimes are equally prosecuted and in which all are punished by death is a social practice in which there is a distinction between criminal and non-criminal conduct, but not one in which there is a continuum of more and less serious crimes. Whatever explains the “ought” of practical rationality, it cannot be analogous to this rigidly dualistic sort of social regulatory system.

5. Trouble in the Philosophy of Mind

It is now possible to show that the meta-theory of practical rationality poses threats to existing research programs in the philosophy of mind. It is not claimed that any of these threats are decisive objections to the theories in question. What I hope to show is simply that philosophers of mind need to grapple with the meta-theory of practical rationality: it matters to other branches of our discipline, and so needs our attention.

The meta-theory of practical rationality meets the philosophy of mind at the theory of the attitudes involved in acting. In the previous section it was argued that whatever explains the “ought” of practical rationality has to stem from the very nature of action itself. If acting is being moved by an intention generated by a belief and a desire, for instance, then the “ought” of practical rationality is something that has to be created by the nature of (some subset of) intention, belief, desire, and their interactions. Notice that it

is vitally important that it be the *nature* of the attitudes involved in action that creates the “ought” of practical rationality. If a theorist explains the “ought” of practical rationality in terms of the common but inessential features of action, then there will be room to imagine actions that lack the inessential features, and so fall outside the domain of practical rationality; there will also be room to suggest that the demands of the “ought” in question are escapable for agents, and so the “ought” in question cannot be the “ought” of practical rationality. Since the nature of the attitudes involved in action is a matter for the philosophy of mind, it is here that there is the most potential for theoretical conflict to develop.

Consider first the neo-Wittgensteinians. In section three it was argued that Sellars and Brandom’s social practices nicely explain certain “ought” facts, but not the ones needed for a meta-theory of practical rationality. But now it seems that this weakness in social practices is more serious, for these social practices are what is essential to the attitudes that make up action, according to these theorists. Being socially treated as committed to certain utterances and movements is what makes it true that I have the relevant attitudes and perform actions, on these views. That is, beliefs, desires, intentions, and actions have little to them beyond the details of how people are subject to particular social practices. If what is essential to these attitudes and their interactions is ultimately a matter of the details of various social practices, then it seems that Sellars and Brandom have no resources for explaining how these attitudes create “ought” facts without appealing to social practices. And social practices have already been shown to be inadequate to the task. The necessity of getting a reasonable meta-theory for practical rationality thus threatens to undermine the whole neo-Wittgensteinian project.

Consider next the teleosemanticists. It was argued that natural selection nicely explains certain “ought” facts, but again not the ones needed for a meta-theory of practical rationality. And again, this weakness is going to have more serious repercussions. The essential features of beliefs, desires, and actions according to Millikan and Papineau are features of their biological functions. Thus, there is nothing to the essential nature of these attitudes beyond the details of their histories of natural selection, and thus, there is nothing but natural selection to appeal to in explaining how these attitudes create the “ought” facts of practical rationality. And natural selection has already been shown to be inadequate to the task. The necessity of getting a reasonable meta-theory of practical rationality thus threatens to undermine the whole teleosemantic project just as much as the neo-Wittgensteinian project.²⁷

Finally, turn to a very different, and (I suspect) substantially more popular theory of mind than either neo-Wittgensteinianism or teleosemantics.

One sort of functionalism holds that to be a belief, desire, intention, or action is to be a node in a causal network of the right sort, where the right sort of network is defined by the platitudes that people hold regarding beliefs,

desires, intentions, and actions. David Lewis has given a particularly precise formulation of this sort of view, but he is far from the only philosopher who has favoured it. On this sort of view, the essence of being an attitude is that of playing (at least most of) the causal roles we commonly associate with the attitude. But it is far from obvious that the causal roles we most commonly associate with the attitudes are causal roles that are conceptually connected to evaluative standards. In fact, it would seem that the causal roles we associate with beliefs, desires, and intentions have nothing at all to do with evaluation, and so are not in a position to explain the creation of any evaluative standard at all. The causal roles associated with desire, for instance, are things such as causing actions that might bring about what is desired, causing feelings of pleasure when what is desired is obtained, causing agents to think longingly of what is desired, and so on. Mere causal interactions of this generic sort are not sufficient to create “ought” facts. Causal interactions that add up to processes of regulation or design suffice to create “ought” facts, as was shown earlier. But there is nothing in our commonsensical understanding of the causal roles of desires that suggest regulation or design. We do not generally think of our desires as regulating or designing our actions. If they do, in some sense, regulate or design our actions this is not part of our commonsensical knowledge of them.²⁸ Hence this familiar sort of functionalism contains nothing in its theory of the attitudes that might allow for the “ought” of practical rationality to be explained. A different sort of functionalism, one that made the regulation or design of action part of the essential nature of desire (or of evaluative belief – there are various options here), might well be in a position to explain the “ought” of practical rationality. But commonsense functionalism is ill-suited to the task.

The three sorts of theories just considered far from exhaust the theories of mind that exist. And the arguments against these three sorts of theories are far from decisive. As stated at the outset, my purpose in this section has not been to show that certain theories are wrong, but to show that there is an important set of considerations that must be addressed by philosophers of mind when building their theories of the attitudes (at least), and that philosophers of mind who neglect these considerations do so at their peril. If neo-Wittgensteinians, teleosemanticists, and commonsense causal-role functionalists all face serious problems explaining how instantiating PR adds up to being rational, the problem is probably one all theorists of the attitudes should address.²⁹

6. On from Here

The theory of the attitudes is a central part of the philosophy of mind. For this reason, any new insight into the attitudes is one that has

ramifications for the field. If practical rationality is a real phenomenon, warranting a substantive meta-theory, then there are ramifications. Any adequate substantive meta-theory of practical rationality will have to explain how “is” facts about the natures of the attitudes involved in acting add up to the “ought” facts of practical rationality. And so any adequate theory of the attitudes will have to permit such an explanation. Whether this means a theory of the attitudes that involves regulation, design, or something else, it is something that is not obviously compatible with existing theories of the attitudes. It is something that is going to take work. Will the same work be required to accommodate a substantive meta-theory of epistemic rationality? That, too, remains to be seen.

Notes

1. I suppress all subtlety here for the sake of exposition.
2. ‘Consistency’ might seem an evaluative term, but it is not. See section two for a discussion.
3. Or they are theories that explain away the apparent existence of the “ought” facts. See section two.
4. I use the expression ‘adds up to’ here to be noncommittal about the precise details of the relation between “is” facts and “ought” facts. So far as I am concerned, a paradigm of X facts adding up to Y facts is found in chemistry: H₂O facts add up to water facts. But since how to understand this relation is disputed even for chemistry, I eschew any more specific theory here. The reader’s preferred way of understanding the relation in chemistry should serve equally well for understanding the counterpart relations in morality and practical rationality.
5. Others who have proposed theories, or constraints on future theories, of practical rationality include Quinn (1990) and Bratman (1999).
6. Perhaps most prominent at present is Scanlon’s non-reductive account which gives a theory of practical rationality in terms of practical reasons, but which treats the property of being a reason as irreducible. See Scanlon (1998).
7. Gibbard (1990).
8. The word ‘rational’ is thus akin to words like ‘yummy’, which are appropriate when one approves of a thing but which are held to be principally used to express feelings rather than to describe reality.
9. For a very recent critical work, see Schroeder (2008).
10. Stich (1990).
11. See, e.g., Kripke (1972), Chalmers (1996). For a version of this complaint more sympathetic to materialism, see, e.g., Levine (1983).
12. Some hold out hope of explaining why this lack of explanation is not decisive against materialism (e.g., Hill 1997 among many others). These arguments typically rely on special facts about the ways in which we conceive of consciousness, and so are not relevant to arguments about practical rationality.
13. Or making probable, or typical—set aside these details.
14. At least, this is one reasonable theory of self-defeatingness. Other (e.g., more Kantian?) theories are possible, but appear to lead to similar conclusions.

15. This is discussed in more detail in Schroeder (2003).
16. Sellars (1963), Brandom (1994).
17. Sellars seems to give a reductive account. Not so for Brandom, though Brandom does not disallow the possibility of a reductive account.
18. See Schroeder (2004a) for further discussion.
19. I borrow the example from Arpaly (2003).
20. The best case for this is still found in Millikan (1984); refinements and a survey of the main competition can be found in Godfrey-Smith (1993; 1994). For important criticisms, see Matthen (1997), Davies (2001).
21. See, e.g., Dennett (1995).
22. More on the conceptual necessity of this in Schroeder (2004a).
23. The example alluded to appears in Davidson (1987).
24. Note the tension between this claim and Hursthouse's idea that there are "arational" actions. See Hursthouse (1991).
25. Davidson's theory of action is a well known example. See Davidson (1980), chapter 1.
26. Clearly articulated in, e.g., Korsgaard (1996).
27. Among the teleosemantic works, Dretske (1988) presents a particularly interesting case. In this work, Dretske derives his functions from natural selection over neural populations (roughly) as an organism learns, rather than over populations of whole organisms as the species evolves. This makes it much more difficult to accuse his view of the problems facing, e.g., Millikan and Papineau.
28. It might, however, be a result that follows from the theory of desire in Schroeder (2004b), if it is true that causing reward learning in a system amounts to regulating that system.
29. Davidsonians should not be confident that they escape the problem either, in spite of the common claim that a Davidsonian theory of the mind captures the normativity of the mind. See Schroeder (2003) for a critique.

References

- Arpaly, N. 2003. *Unprincipled Virtue*. New York: Oxford University Press.
- Brandom, R. 1994. *Making It Explicit*. Cambridge, MA: Harvard University Press.
- Bratman, M. 1999. Toxin, Temptation, and the Stability of Intention. In *Faces of Intention: Selected essays on intention and agency*. New York: Cambridge University Press. 58–92.
- Chalmers, D. 1996. *The Conscious Mind: In search of a fundamental theory*. New York: Oxford University Press.
- Davidson, D. 1980. *Essays on Actions and Events*. Oxford: Clarendon.
- Davidson, D. 1987. On Knowing One's Own Mind. *The Proceedings and Addresses of the American Philosophical Association* 60, 441–58.
- Davies, P. 2001. *Norms of Nature: Naturalism and the nature of functions*. Cambridge, MA: MIT Press.
- Dennett, D. 1995. *Darwin's Dangerous Idea: Evolution and the meanings of life*. New York: Simon and Schuster.
- Dretske, F. 1988. *Explaining Behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.

- Gibbard, A. 1990. *Wise Choices, Apt Feelings: A theory of normative judgment*. Cambridge, MA: Harvard University Press.
- Godfrey-Smith, P. 1993. Functions: Consensus Without Unity. *Pacific Philosophical Quarterly* 74, 196–208.
- Godfrey-Smith, P. 1994. A Modern History Theory of Functions. *Nous* 28, 344–62.
- Hill, C. 1997. Imaginability, Conceivability, Possibility, and the Mind-Body Problem. *Philosophical Studies* 87, 61–85.
- Hursthouse, R. 1991. Arational Actions. *Journal of Philosophy* 88, 57–68.
- Korsgaard, C. 1996. *Sources of Normativity*. New York: Cambridge University Press.
- Kripke, S. 1972. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Levine, J. 1983. Materialism and Qualia: The Explanatory Gap. *Pacific Philosophical Quarterly* 64, 354–361.
- Lewis, D. 1972. Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy* 50, 249–58.
- Matthen, M. 1997. Teleology and the Product Analogy. *Australasian Journal of Philosophy* 75, 21–37.
- Millikan, R. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Papineau, D. 1987. *Representation and Reality*. New York: Basil Blackwell.
- Quinn, W. 1990. The Puzzle of the Self-Torturer. *Philosophical Studies* 59, 79–90.
- Schroeder, M. 2008. *Being For: Evaluating the semantic program of expressivism*. New York: Oxford University Press.
- Schroeder, T. 2003. Donald Davidson's Theory of Mind is Non-Normative. *Philosophers' Imprint* 3, 1–14.
- Schroeder, T. 2004a. "Functions from Regulation." *The Monist* 87, 115–35.
- Schroeder, T. 2004b. *Three Faces of Desire*. New York: Oxford University Press.
- Sellars, W. 1963. Some Reflections on Language Games. In *Science, Perception, and Reality*. London: Routledge and Kegan Paul. 321–58.
- Scanlon, T. 1998. *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- Stich, S. 1990. *The Fragmentation of Reason*. Cambridge, MA: MIT Press.