Communist Conventions for Deductive Reasoning

Sinan Dogramaci
The University of Texas at Austin

Abstract

In section 1, I develop epistemic communism, my view of the function of epistemically evaluative terms such as 'rational'. The function is to support the coordination of our belief-forming rules, which in turn supports the reliable acquisition of beliefs through testimony. This view is motivated by the existence of valid inferences that we hesitate to call rational. I defend the view against the worry that it fails to account for a function of evaluations within first-personal deliberation. In the rest of the paper, I then argue, on the basis of epistemic communism, for a view about rationality itself. I set up the argument in section 2 by saying what a theory of rational deduction is supposed to do. I claim that such a theory would provide a necessary, sufficient, and explanatorily unifying condition for being a rational rule for inferring deductive consequences. I argue in section 3 that, given epistemic communism and the conventionality that it entails, there is no such theory. Nothing explains why certain rules for deductive reasoning are rational.

Introduction

Epistemologists continue to search for the true theory of rationality. In just the last few decades, many important theories have been developed, theories with names like 'reliabilism', 'conservatism', 'the metasemantic account', and 'the pragmatic account'.[1] More and more, the playing field is coming to look like it has for some time in ethics, where the familiar competitors include utilitarianism, Kant's categorical imperative, ideal observer theories, and contractualist theories.

Of course, as in ethics, there is nothing approaching a consensus among epistemologists. Every theory so far aired is known to face objections and remains highly controversial. But, even if no one has yet published the correct theory in all its details, epistemologists remain optimistic, showing few signs of calling off the search for the true theory.

I am worried that this optimism is a mistake. My worry derives from a view I have about the *function* (in the sense of the point, or purpose) of our epistemically evaluative language. I first introduced this view in Dogramaci (2012). Here I will give the view some further elaboration and defense, and then I will argue for a conditional: if this view of the *language* is right, then it follows that there is no theory of epistemic rationality, that is, no theory of the *property* of being rational.

For the sake of keeping discussion focused and manageable, this paper restricts itself to the topic of *deduction*, by which I shall just mean inferences to consequences

that logically follow from one's premises. (I do favor generalizing things to cover non-deductive cases; again see Dogramaci (2012).) Given the restriction to deduction, the paper will consequently focus more on theories that are tailored to explaining rational deduction, in particular metasemantic and pragmatic approaches, rather than reliabilist and conservatist approaches.

The paper has three sections to follow. Section 1 develops my view of our epistemically evaluative language. Section 2 examines proposals offered by traditional epistemologists seeking a theory of rationality, and from these proposals extrapolates a general model of what a theory of rationality is supposed to look like. Section 3 then argues, on the basis of a special kind of conventionality I find in the meaning of epistemic evaluations, for my view that there is no such theory of rational deduction.

1. Epistemic Communism

Here, in section 1, I'll motivate (1.1) and introduce (1.2) my view of the function of epistemically evaluative language, and I'll address some potential misunderstandings (1.2 – 1.3) and worries (1.4). I call the view *epistemic communism*.

1.1 Motivation: Some Deductions are Valid but Not Rational

Take the old story, which everyone already knows, of Fermat writing in the margin of his copy of the ancient Greek math text, Diophantis's *Arithmetica*, that he'd proved the conjecture today known as Fermat's Last Theorem. We don't know what Fermat's argument was, but imagine it was very simple: imagine his argument began with some standard arithmetical axioms, and then, with no intermediate steps, no intermediate lemmas, he took an inferential leap right to the theorem. Would you say his belief in the theorem was rational? Of course not, not on this story. Even though we now know, centuries later, that such a transition is *valid*, it is not a transition that can confer *rationality* on the conclusion. (Rationality, justification, reasonableness: same thing. Use whichever word you like.)

Let's call the *epistemic rule*—that is, the *belief-forming method* or *process*—for inferring Fermat's Last Theorem just from arithmetical axioms *Fermat's Rule*. And when an epistemic rule is capable of conferring epistemic rationality on its conclusion, let's call it a *rational rule*. (I say 'capable' because even a good rule may not always produce a rational conclusion, for example, if it is applied using basing beliefs that are not themselves rational.) So, we can say that Fermat's Rule, though valid, isn't rational.

It's easy to multiply examples of rules that are valid but not rational. Take any unobvious logical truth. Take $((p \supset q) \supset p) \supset p$, known as Peirce's Law. Let's use *Peirce's Rule* to label an epistemic rule for inferring Peirce's Law immediately, from no premises. Then, Peirce's Rule is valid, but not rational. There is actually some debate over how Fermat's Last Theorem can be logically derived from standard arithmetical axioms (see, e.g., McLarty (2010)), so from here on, let me use Peirce's Rule for my main concrete example of a valid but not rational rule.

Only a very small handful of the infinitely many valid rules are rational rules. What are some rational valid rules? One is *Modus Ponens*, a rule for inferring *q* from the premises *p* and $p \supset q$. I confess, the details of the content of the epistemic rule here are actually very hard to fill in. For your belief in *q*, inferred from these premises, to be rational, there is a further condition: you need to lack any sufficiently strong prior reasons for doubting *q*. When there are such reasons to doubt *q*, the rational rules would have you give up belief in either *p* or $p \supset q$. It's very hard to say exactly when one has such a sufficiently strong prior reason for doubting *q*, and so it's hard to fill in the details of the rule. But, details aside, I think it's clear enough that we do follow some rule roughly along these lines, and it's one of the few valid rules that's rational. (Note I'm not disagreeing with Harman (1986), who allows there's a rule of this sort, but would call it a rule of 'immediate' implication, and would urge that it reveals no *special* relationship between an epistemic rule's being rational and its being logically valid.)

Other rational rules are recognized by familiar names: *Conditional Proof*, *Conjunction Introduction*, *Conjunction Elimination*, *Addition* (a.k.a. *Disjunction Introduction*), *Reasoning by Cases* (a.k.a. *Disjunction Elimination*), *Reductio ad Absurdum*, and a small handful of others. For all of these rules, it's again hard to specify the exact content, but they're roughly well enough understood for our purposes in this paper. I'll focus on Modus Ponens from here on.

In this paper, I am concerned only with *basic* rules, rules we follow not by following other rules. You might decide to start following a non-basic epistemic rule for inferring Peirce's Law; this could be a rule you follow by following another rule for inferring whatever you've seen proved. Peirce's Rule, however, is a basic epistemic rule.

I'm saying Modus Ponens, and not Peirce's Rule, is among the rational rules in the sense that it is capable of leading to a rational conclusion in deductive reasoning. Some may hesitate to agree. But, I allow there may well be a variety of senses of 'rational' on which Peirce's Rule is rational, even some sense associated with what one ought to believe, and even a sense some might describe as the rule's being capable of leading to a known conclusion. Certainly some philosophers will say of Peirce's Rule that it is rational, maybe as long as the reasoner has, say, a certain accompanying intuition, or a deep commitment to the rule's reliability (see Huemer (2001, 2007) and Foley (1987, 1993) for such views). My claim is only that there is *an* ordinary type of epistemic evaluation, a type we typically use ordinary words like 'rational', 'justified', 'reasonable' and even the ordinary word 'logical' to express, and on this standard of evaluation conclusions reached by Peirce's Rule are not capable of being rational in the way Modus Ponens conclusions are, no matter the reasoner's intuitions or commitments. Plausibly, as Alston (2005) discusses, there are many dimensions of epistemic evaluation. One of these says Modus Ponens is special, it is rational, while Peirce's Rule is not. That's what the sort of evaluation I'm concerned with says. Take it as a premise of the discussion in this paper: Modus Ponens is special; it is rational in a way that other valid rules, such as Peirce's Rule, are not. (This premise is common to the contemporary theories of rational deduction that we'll be looking at later on; see p.14.)

The sort of evaluation I'm concerned with is one that diverges from evaluations of a rule's *reliability*, and hence from evaluations of its *validity*. There is thus an analogy between my evaluation of Peirce's Rule and the evaluations critics use to show that reliabilism doesn't reveal the truth about rationality, or at least not the whole truth. Reliabilism, in its simplest form, says a rational rule just is a reliable rule; see Goldman (1979/92, 1986).) Critics of reliabilism point out that we give a negative evaluation of Norman, the famous unwittingly reliable clairvoyant (see BonJour (1980)). I'm concerned with an analogous negative evaluation of Peirce's Rule and Fermat's Rule (see Boghossian (2003) for an explicit such evaluation of Fermat's Rule). As mentioned, some, like Huemer and Foley, will say of Peirce's Rule, just as some say of Norman's reasoning, that it is rational. Perhaps, understood in a certain way, such evaluations aren't wrong, but they are not the sort of evaluations I'm interested in here. I'm not interested in evaluations that track, as on Huemer and Foley's views, the subject's own intuitions or commitments. The evaluations I'm interested in here are ones that track the subject's use of Modus Ponens, Conditional Proof, Reductio, and the other familiar rational rules.

1.2 Epistemic Communism: The View

Why would any useful or sensible standard of epistemic evaluation privilege Modus Ponens over Peirce's Rule? Of course, it would be psychologically difficult for a real person to follow Peirce's Rule, but my question here arises just because a hypothetical follower of Peirce's Rule is evaluated as drawing an irrational conclusion (again, at least according to that one common, ordinary standard).

The question is also well motivated, because it is puzzling why we should make epistemic evaluations that don't promote true belief in the direct and simple way that a reliabilist's evaluations do. (I elaborate why this is so puzzling at greater length in Dogramaci (2012).)

I offer the following explanation for why we evaluate the follower of Peirce's Rule—call him 'Peirce' for short—as drawing an irrational conclusion. Imagine we put together a team consisting of ordinary mathematicians plus Peirce, and the team's fundamental goal is to maximize the number and interestingness of the results they prove while minimizing on costs of time and intellectual resources. Peirce immediately announces he's proved that $((p \supset q) \supset p) \supset p$. Since this is in fact a valid theorem, he might seem to be doing well at trading off results against effort. But, what will the other team members say about Peirce's proposed result? Since they don't follow Peirce's Rule, they will not accept Peirce's (one-line) proof of his result. Although Peirce's proof is in fact valid, the others need some way of checking this for themselves, some way of checking that Peirce isn't making wildly unreliable claims. Unless the team somehow checks up on reports like this, they are at risk of gullibly buying testimony from anyone at all. Peirce's result should only be acceptable to the others once a proof has been offered which only uses rules they accept, rules like Modus Ponens. But, if team-wide acceptance of the result has to wait for such a "slow" proof, then Peirce's work offers no advantage as far as the team's progress goes. Insofar as team-progress is our concern, it thus makes sense

to negatively evaluate a conclusion reached using Peirce's Rule. The team's progress proceeds most efficiently when everyone uses the same rules. There is thus a social epistemic virtue of coordination among rules.

(As always, I mean here our basic epistemic rules, rules for reasoning, not rules for producing derivations or proofs. Realistically, a mathematician could of course write down Peirce's Law as a step in a fine proof without supplying further support. Thanks to a referee for suggesting this important clarification.)

Our actual negative evaluation of Peirce's Rule and our positive evaluation of Modus Ponens (and other rules) reflects a social, epistemic benefit gained from coordination of epistemic rules. The benefit is the efficient division of labor among members of a group aiming to discover and share interesting results with a minimum of cost. A division of labor is efficient here because the product of each worker's labor, their discoveries, can be safely spread to others. When we all participate in a practice that promotes coordination among our epistemic rules, we can all safely trust the *testimony* of others. Reports of your discoveries then become as valuable to me as my own discoveries, and, likewise, my reports are valuable to you. With uncoordinated outsiders, like Peirce, it is either dangerous or useless to accept their testimony, because we either cannot confirm their reports, or we have made our own discovery already. But, in a community where we work to coordinate epistemic rules, spreading beliefs through testimony becomes both safe and valuable. I call this account of the function of our epistemically evaluative practice *epistemic communism*.

Although this idea of epistemic communism is a new proposal (I first proposed it in Dogramaci (2012)), I think that once you hear it, it's fairly intuitive. The idea is just that, by sharing our epistemic rules, we each become the other's *epistemic surrogate*. A few epistemologists have already broached this issue of the function of epistemic terms: in particular, Craig (1990) and Reynolds (2002) offered an insightful proposal about another epistemic term, the term 'knows'. Their proposal, in Craig's words, is this: 'to put it briefly and roughly, the concept of knowledge is used to flag approved sources of information' (p.11). They were on to the idea that we want to rely on others as our epistemic surrogates, but what they left out is the communist's idea that we achieve this via coordination. Coordination is an especially efficient way of pursuing the truth. In effect, it's the efficiency of parallel processing. We are a community of *parallel epistemic processors*.

(Gibbard (1990) argued that the function of normative discourse is coordination, but his view of the benefits of coordination is unrelated to the communist's proposal about testimony; also, he focused on practical, rather than epistemic, evaluation.)

Let me add six clarificatory comments about what epistemic communism says, and more importantly what it does not say.

First, I do not claim that coordination is by itself *sufficient* for successfully discovering lots of interesting theorems. We must coordinate upon valid rules. A coordinated team that follows an invalid rule will of course do poorly. The team's stock of rules should be sound (valid), and hopefully also complete. Other than soundness and completeness, is there anything else necessary for a team's success? You might think that, since the team is trying to maximize certain outputs while

minimizing certain costs, they do better by including more and more rules that allow for proofs of more results in fewer steps. But, each additional accepted rule incurs its own cost on resources, so it's not clear there's any benefit to adding all that many rules to an already complete set.

Second, I also do not claim that coordination upon a stock of sound and complete rules is *necessary* for discovering lots of theorems. Suppose Peirce follows not only Peirce's Rule but lots of other valid yet irrational rules and, after a while, the team confirms that Peirce has a highly reliable track record. Then we could trust his reports even without sharing his same (basic) epistemic rules. Trust without coordination is possible. However, it's not an *efficient* way of pursuing our goal (of discovering interesting truths). It requires the costly work of independently checking track records. The more efficient way to safely trust testimony is by coordination. Note also that there need not be any downside to suppressing Peirce's exotic rules: we can advise Peirce to believe what he finds himself disposed to believe, but armed now with the knowledge that there is a reliable track record for what he is disposed to believe. Then we will all be coordinated in our basic rules, one of which is a rule to trust anyone and anything with a reliable track record.

Third, epistemic communism presupposes that epistemic evaluations influence others' behavior, including aspects of how they reason, but it doesn't implausibly presuppose any metaphysically necessary connection here. The connection may be an innate, evolved feature of human psychology. Darwin (1871/74, chapter 5): '[A]nother and much more powerful stimulus to the development of the social virtues [than self-interested calculation], is afforded by the praise and the blame of our fellow-men. To the instinct of sympathy, as we have already seen, it is primarily due, that we habitually bestow both praise and blame on others, whilst we love the former and dread the latter when applied to ourselves; and this instinct no doubt was originally acquired, like all the other social instincts, through natural selection.'

Fourth, I don't claim that epistemic evaluation is the only way that coordination is fostered; evaluation is just one way. Plausibly, we have an evolved, innate disposition to accept a certain coordinated stock of epistemic rules. Epistemic evaluation is a supplementary tool for policing others, for keeping them on the straight and narrow path of following exactly the coordinated rules. It's an especially useful tool for a species, like ours, in which individuals are much better at catching others' mistakes than at catching their own (see Pronin et al. (2002)).

Fifth, it's important to note that you don't need examples as far-fetched as Peirce or Fermat to illustrate failure of coordination. Coordination can fail in two broad ways: failures in competence and failures in performance. (The distinction here is, of course, the analog, for epistemic rules, of Chomsky's famous competence/performance distinction for rules of syntax.) In the example above, Peirce fails in competence: he accepts a rule that we don't want him to follow, and his inference manifests this acceptance. More ordinary examples of competence failure might be reasoning that follows some rule of wishful-thinking, fear-based thinking, or prejudicial thinking. In the other sort of failure, a reasoner might accept only

approved rules, but performance failures prevent success at following those rules. So, for example, Smith might accept only approved rules, but immediately infer an instance of Peirce's Law because he somehow—who knows exactly how epistemic rules are cognitively processed?—wrongly takes it to be an instance of the approved rule licensing immediate belief in conclusions of the form (A ⊃ A). Our epistemic evaluations promote coordination by suppressing both of these ways in which coordination can fail, failures in competence and failures in performance.

Sixth, epistemic communism makes no explicit commitment to one theory or another of the conditions for rationally believing testimony. It is naturally paired, however, with the reductionist view of Hume (1748/2000), Fricker (1994), Adler (2002, chapter 5), and Lackey (2008, chapters 6–7). Reductionists say that, to rationally believe testimony, you need empirical reasons to believe the testimony reliable, and these reasons must be ultimately independent of your reliance on testimony. Anti-reductionists, such as Reid (1780/1981), Coady (1992) and Burge (1993), say you can rationally trust testimony without such independent empirical confirmation of its reliability. (Some authors use these labels in slightly different ways. See Lackey (2008, p.195) for why she is a reductionist on my labeling, but not her own.) Reductionism enjoys much initial plausibility from the intuitive thought that we should not gullibly accept others' testimony; we should somehow protect ourselves against buying false testimony. A main argument for anti-reductionism is that it would seem to be an impossible task to check up on all the sorts of things we learn by testimony, and thus reductionism would threaten skepticism. But, this argument for anti-reductionism ignores the communist's very practicable way of confirming that testimony is reliable: actively police others to enforce that they follow the same epistemic rules you follow. When others share your rules, their observations and reasoning are then as good as your own observations and reasoning.

(Consumers of testimony still have work to do before they can know others' *testimony* is reliable, since reliable testimony requires not only reliable belief but also honest reporting. Communism helps with the first condition; other research has addressed the second. On this second issue, a recent survey examining the mechanisms humans have for effectively confirming honesty in others is Sperber, Clément, Heintz, Mascaro, Mercier, Origgi and Wilson (2010).)

Epistemic communism is a view about the pragmatics of epistemic evaluation. The question it directly addresses is one about why we do, and why we should, make certain evaluations, such as our evaluations that an inference by Modus Ponens is rational and one by Peirce's Rule is not. It answers that question with a view about the role and utility in our social lives of that linguistic practice. But, does it provide any answer to the questions: why is Modus Ponens epistemically rational, and why is Peirce's Rule irrational? I want to show how epistemic communism does shed light on those questions (by showing they have no answer). But before I get to that, sub-section 1.3 clarifies the relationship between communism and certain views concerning the semantics of epistemic evaluations, and sub-section 1.4 responds to an interesting and important worry one might have about epistemic communism.

1.3 What Semantics Best Suits the Communist's Pragmatics?

It may seem to some that the examples that motivate communism also support a subject-sensitive or contextualist semantics for epistemically evaluative terms. In this sub-section, I argue there is no such support here for those semantic views.

I have focused on an ordinary use of 'rational' according to which Modus Ponens is rational and Peirce's Rule is not, regardless of what the reasoner's intuitions or deep commitments may be. While other uses of 'rational' may express a more subjective notion, the notion I am focusing on is one where, I claim, we speak truly if we say, 'Modus Ponens is rational, and Peirce's Rule is not, *necessarily*; the facts about which epistemic rule a reasoner uses settle whether her rules are rational, no matter the further details of the reasoner's personal or social context.' Some readers, even those most sympathetic to communism and its motivating examples, may be tempted to dispute this. They may be tempted to claim the above evaluation is false because the following evaluation is true and inconsistent with the above one: 'If a reasoner were a member of an alien community where everyone else has coordinated upon a stock of sound and complete rules for deductive reasoning that includes Peirce's Rule and excludes Modus Ponens, then this reasoner would form rational beliefs by using Peirce's Rule, and not by using Modus Ponens.' If that is right, it would suggest epistemic evaluations are subject-sensitive, sensitive, in particular, to whether a reasoner's rules are coordinated with those rules used throughout her own local community.[2]

I appreciate that there is some temptation, even for a communist sympathizer, to say 'Modus Ponens is possibly irrational; in particular, it is irrational when used by one person in the described alien community where no one else uses it'. But, that temptation is due to a natural confusion, and does not favor a subject-sensitive semantics. We feel a temptation to say this because, as soon as we consider the alien community, we easily recognize that the aliens best serve the coordinative function of epistemic evaluation by making evaluations that reinforce their already common use of Peirce's Rule and suppress any occasional deviant uses of Modus Ponens. The aliens could serve the coordinative function by speaking a language homonymic to English, though their use of 'Modus Ponens is rational' expresses a falsehood; their evaluative term 'rational' has a different extension from our term. It's plausible this is how they would speak, and speak truly, if they were as similar to us as could be while having different habits of reasoning (it's overall more similar to us to use 'rational' to coordinate, than to use it to promote what is, for them, a deviant rule). It is thus easy for us to confuse their term 'rational' with our homonymic term 'rational', and to lapse into using their term without realizing it. Perhaps this flexibility, letting ourselves easily slide into speaking a subject's language, serves some function; maybe it's useful for trying to better understand foreign subjects and their cultures. But, it does not serve the communist's *coordinative* function for *us* to ever say 'Modus Ponens could be irrational, and Peirce's Rule could be rational.' Our reasoning is coordinated upon Modus Ponens and not Peirce's Rule—always has been, and, realistically, always will be. It thus benefits us to ensure that others share our rules, for the sake of efficiently making testimony safe. *We* best serve the coordinative function by adamantly affirming, 'Modus Ponens is rational, always

was, always will be, everywhere and for everyone, no matter what.' The modal force of such evaluative statements only strengthens their power to promote reasoning by Modus Ponens. We want others to use Modus Ponens not just in the instance for which we are praising or criticizing them, but in the widest range of circumstances. This is because we need to plan and prepare for possible circumstances before we know which is actual. (See Craig (1990, pp.19–20), where a similar point is made.)

Why wouldn't we allow exceptions, saying, 'It would be permissible for someone to not use Modus Ponens, as long as it's known by everyone with certainty that that person will never deliver testimony to us.'? Well, why *would* we make exceptions? By communist lights, there is no reason to. It introduces complexity into our evaluations *to no benefit*. We do better to keep our evaluations—our instructions to feeble-minded testifiers—simple. Since it's likeliest that we are (evolutionarily) engineered to take the simplest, most efficient, solution to a problem, other costs and benefits being equal, we likely do not allow such excuses.

There is thus no reason here to think that the communist's view of the pragmatics of our term 'rational' favors a subject-sensitive semantics. What about a contextualist semantics, though, one where utterances of 'Modus Ponens is rational' can, without changing meaning, change truth-value depending on the context of utterance?[3] Again, I see nothing in the communist's pragmatic story that favors any such revisionary view of the semantics of 'rational'. The simplest semantic account, it seems to me, has it that the aliens and we use different homonymic words with different fixed extensions. I have no positive argument against the contextualist view that says the aliens and we use a single word that changes its extension depending on the speaker's community, or depending on what the speaker is considering, or what is salient to her, at the time of her utterance. But these views posit extra semantic complexities that don't serve us any better, as far as communist pragmatic purposes go, than a simple account on which our term, 'rational', has a fixed extension.

1.4 Reply to a Worry that Communism Tells an Incomplete Story

Some may worry that epistemic communism gives a crucially incomplete account of the function of epistemically evaluative discourse. One way this worry can get traction is from the thought that we use epistemic evaluations in contexts where our evaluations could not possibly promote interpersonal coordination, in particular in the context of *first-personal deliberation*. In response to this worry, an initial point is that, even if epistemic evaluations did serve a further non-coordinative function, that of course would not show the coordinative function doesn't exist. Still, it would diminish the importance of the communist proposal. In this section, I will say more in order to directly criticize the idea that there is any further function served by first-personal epistemic evaluations.

But, before turning those criticisms, I'll take a paragraph to acknowledge that some may certainly have worries about whether there are *third*-personal epistemic evaluations that serve some function other than that proposed by epistemic communism. One might worry: do epistemic evaluations serve the coordinative function when embedded in logically complex constructions, when used in causal

explanations or predictions, or when embedded in modal operators? But, I think the worry here is much less urgent than the worry about first-personal evaluations. It's not so hard to see how third-personal evaluations could promote interpersonal co-ordination, even within complex embeddings. Consider an evaluation buried inside a causal explanation: vos Savant switched doors because she is rational. That state-ment not only explains what vos Savant did, but endorses the rules she followed. Or consider an evaluation embedded in a counterfactual: reasoning by Modus Ponens *would be* rational, no matter the contents of p and $p \supset q$, no matter what day of the week it is, no matter whether I am around to watch over you or not, indeed, plausibly, no matter what. So, I hope it is not so implausible that third-personal evaluations do serve the coordinative function, even when the evaluations are not simple cases of calling a particular belief rational (or irrational).

Let's turn, then, to the worry about first-personal evaluations. The worry is that, since first-personal evaluations cannot promote interpersonal coordination, they do not serve the coordinative function. To be sure, a first-personal eval-uation might still serve as a useful self-criticism to help ensure that my future reasoning is reliable (reliable by my present lights, at least). That is still a coor-dinative use of evaluation, only aimed at *intra*personal coordination. The worry is that first-personal evaluations serve some very different function by playing some special role other than the coordinative role. If that were so, the commu-nist's story would be missing something important. So, let's see if there is any such thing.

What special role might it be thought that epistemic evaluations play in first-personal deliberation? The natural idea is that epistemic evaluations do, can or should play some sort of *guiding* role in good, proper, or serious deliberation. My impression, for whatever it's worth, is that this vague suggestion (that first-personal evaluations guide deliberation) finds widespread sympathy among philosophers, though, perhaps due to its vagueness, it is not quite as widely articulated in print. Still, the idea finds its way into current work. Here is Roger White, for one, endors-ing the idea: 'Justification is a kind of guide to the truth. We seek to form justified beliefs as a means to forming true beliefs. This is why a rational inquirer is sensitive to questions of justificatory status in forming his beliefs. In a serious inquiry as to whether P, we ask ourselves whether we would be justified in believing P.' White (2006, p.539) (The idea that the concept of rationality guides proper deliberation can also be found in BonJour (1985, pp.7–8) and Pollock and Cruz (1999, p.14; also chapter 5), though they endorse the idea in more cautious, and more vague, terms than White's.) If it were indeed true that 'serious inquiry', as White says, involves the thinker's asking herself whether she would be justified (or rational) in believing this or that, then that would have significant implications. It would at least strongly suggest that, in deliberation, the concept of rationality is serving some exclusively first-personal function.

It seems to me, however, that White's description gets the structure of deliberation backward. Consider the process by which a thinker forms beliefs. Most beliefs are formed spontaneously. It's obvious that the process by which we spontaneously form beliefs does not make any essential use of the concept of rationality.

Now, the label 'deliberation', though, is usually reserved to describe a more careful, considered or reflective process of belief-formation (though such belief-formation is probably still non-voluntary). But still, contrary to the popular view that White expresses, even the most considered belief-forming process does not essentially involve the use of the concept of rationality. A thinker *may* use the concept of rationality in the course of her deliberation, but it would be totally inessential to the deliberative process. Imagine deliberating about something, say, whether there is alien life. How do you deliberate about whether there is alien life? You ask what the *data* is: what do we see through telescopes, what do radio signals from space indicate, how many other planets are as hospitable as ours? You ask whether that data provides a *reliable* indication of the existence of alien life. By answering those questions, you may settle your deliberations about whether there is alien life. Do you, or should you, also ask whether the *rational* thing to believe, given that data, is that there is alien life? No, that is not normally relevant. Or imagine, to return to the deductive side of reasoning, deliberating about whether Peirce's Law, $((p \supset q) \supset p) \supset p$, really is a theorem (as you might have just recently done, if you had not encountered it before reading this paper). You've settled the matter when you've settled whether there are any truth-value assignments that invalidate the claim. Do you also take up the question of whether it is also *rational* to think there is such a truth-value assignment? Again, that is inessential, and not normally relevant. For the most part, an ideal deliberator may proceed perfectly properly without giving any consideration to questions of rationality.

There are exceptional cases. In an American criminal trial, a juror must deliberate to specifically settle whether there is a *reasonable* doubt about the defendant's guilt. Here the normative question of what it would be rational to believe (or to doubt) becomes relevant. Or imagine a bizarre case where you are suffering from, and trying to cure yourself of, epistemic akrasia: your belief about whether p conflicts with your belief about whether it's rational to believe p. Here again, deliberation about p might reasonably be influenced by deliberation about whether it's rational to believe p. But, such exceptional cases only demonstrate that deliberation as to whether p does not generally involve or depend on asking oneself whether it is rational to believe p. The view White expresses thus gets things backward.[4]

It's not at all clear how to make sense of the strangely popular idea that the concept of rationality guides deliberation. Is there some other way, perhaps some much less obvious way, that epistemic evaluations could turn out to play a special role in first-personal thought? I'm unsure how commonplace the following idea is, but maybe some people would find it tempting, so let me quickly address it.

The idea is that we need to make first-personal evaluations in order to fully realize our epistemic potential, and this is because we need to make such evaluations in order to be able to respond to a certain kind of defeater, *undercutting defeaters*. If one were unable to respond to such defeaters, perhaps one would not be irrational, perhaps the defeaters would not exist for you, but you would be an overall less reliable-and-powerful inquirer.

Why think undercutters assign a special role to first-personal epistemic evaluations? Weatherson (manuscript: §2.1) has the idea; he says: 'The picture I'm

suggesting is that undercutting defeaters [for believing H] are in the first instance rebutting defeaters for [note: rebutting defeater for = reason to believe the contrary of] the claim that the agent is justified in believing H. That is, when the agent gets an undercutting defeater, the first thing that happens is that the second-order claim that the agent is justified in believing she is justified in believing H becomes false.' Thus construed, undercutters straightforwardly involve a negative first-personal epistemic evaluation. So, if this is the whole story about how undercutters are generated, then there is a special first-personal role played by epistemic evaluations, one which epistemic communism fails to account for.

Weatherson offers one proposal for how undercutters are generated. A different proposal is that undercutters are generated by evaluations just of the *reliability* of one's belief-forming process, rather than (as Weatherson proposes) an *epistemically normative* evaluation of one's own belief (that it's justified, rational, etc). The two proposals are not inconsistent; there might be two ways to generate an undercutter. I won't claim that Weatherson's proposal is wrong. But, I claim the reliability proposal is right. It is plausible that evaluations of reliability suffice for a reasoner to acquire undercutters. It follows that no use of the concept of rationality is *required* in order for a reasoner to be able to have undercutters. So, there's still no argument I'm aware of showing that, once you have the concept of reliability, the concept of rationality serves any further valuable function in first-personal deliberation. So, I see no reason to think epistemic communism tells an incomplete story about the function of epistemic evaluations.

2. What Is a Theory of Rationality Supposed to Do?

Section 1 argued for a view of epistemically evaluative language, language such as our word 'rational'. Taking that view of the language as a premise, this section and the next will argue for a view of the rational rules themselves. I will argue there is no theory of these rules. This section sets the argument up by explaining what a theory of rationality is supposed to be, and thus what I will be denying exists. The next section, 3, then makes the argument.

A theory of rationality (or justification, same thing) is a kind of *explanation*. The data, the things to be explained, are ultimately just rational beliefs, and the theory *explains why* those beliefs, and not others, are rational. We already know how to take a first step toward explaining why, say, Smith's belief in Peirce's Law is a rational belief: the first step is to observe that Smith formed her belief using only certain epistemic rules, using only Modus Ponens, Conditional Proof, and other such rational rules. But, what's still lacking, and what epistemologists have been struggling to uncover, is the next step, the explanation of why these certain epistemic rules, and not others, are rational. What explains why just these certain basic rules generate and transmit epistemic rationality?

Recent proposed theories of rationality are proposals for how to answer that question. As I see it, the proposed theories all follow a particular model of explanation. Explanation can be different things in different contexts, and philosophers don't all agree on what makes something a genuine explanation, but the

theorists who've proposed theories of rationality seem to be presupposing a shared model. Let me explain what the model of explanation is that I see these theorists adopting, and then I can support my attribution of the model to the proposed theories.

Their model is a *unificationist* model of explanation. A proposal of a theory of rationality is a proposal of a property as *necessary and sufficient* for being a rational rule, and the property is furthermore proposed to *unify* the rational rules. Now, we all understand what a necessary and sufficient property (or condition, same thing) is, but what is it for a property to also *unify* the things it is necessary and sufficient for?

To understand what kind of unification theorists are looking for, consider why theorists of rationality would care about finding something beyond a necessary and sufficient property for rationality. They care because lots of necessary and sufficient properties for being a rational rule are not at all explanatorily illuminating. Just consider the trivial property: being a rational rule. Or consider the highly disjunctive property: being either Modus Ponens or Conditional Proof or Reductio, and so on. A theory of rationality must identify a necessary and sufficient property that is also explanatorily illuminating. I use the word 'unify' for whatever extra feature will make a necessary and sufficient property explanatorily illuminating in the way theorists hope for. (I borrow 'unify' from the unificationist model of explanation some have defended in philosophy of science; see Friedman (1974) and Kitcher (1981, 1989) for classic general presentations of the view that scientific (as well as mathematical) explanation is unification.)

The epistemologists themselves do not usually use the word 'unify' for this extra feature, but they gesture at the same idea with similar language. Let's look at what they have said.

Goldman (1979/92), when he introduced us to reliabilism, prefaced the proposal with a number of insightful remarks about what a theory of rationality is in general supposed to be. (Again, reliabilism is the proposal that a rule is rational if and only if, and because, it is reliable. The proposal was not tailored to address the case of rules for deductive reasoning, and so it is particularly implausible for explaining why Modus Ponens, rather than other valid rules (and thus equally reliable rules) like Peirce's Rule, is rational.) Goldman said: 'I want a set of substantive conditions that specify when a belief is justified.' (p.105); 'Since I seek an explanatory theory, i.e., one that clarifies the underlying source of justificational status, it is not enough for a theory to state "correct" necessary and sufficient conditions. Its conditions must also be appropriately deep or revelatory.' (p.106); and 'A theory of justified belief of the kind I seek ... must be couched at a suitably deep, general, or abstract level' (p.106).

Boghossian (2003) also prefaced his more recent proposal for a theory of rational deduction with language suggesting a similar idea. He said he is seeking a property *in virtue of which* a rule for deductive reasoning is rational (he writes in terms of facts rather than properties, but it makes no difference): 'I am asking *by virtue of what facts* a deductive inference transfers warrant, and not just under what conditions it does so' (p.226).

Enoch and Schechter (2008) again gesture at the same idea; they echo Goldman's use of 'substantive', and they are explicit about seeking an explanatorily illuminating condition common to just the rational epistemic rules: 'It is highly implausible that it is merely a brute fact that we are justified in employing certain methods as basic and not others. It is much more plausible that there is a principled distinction between the two classes. And it is plausible that this distinction is one that we would be happy to accept as relevant to justification, one that presents the relevant methods in a rationally positive light. This, then, is what a theorist should be after: a substantive account of what's in common to the belief-forming methods we are justified in employing as basic, an account that provides a plausible explanation of this epistemic status.' (pp.548–9)

Horwich (2008) surveys a number of mainstream proposed theories of rationality and explicitly says (see pp.456–7 and footnote 9) that (on the most charitable interpretation) they are aiming to explain the rationality of our rules by finding a single, overarching principle—an 'ur-norm' as he calls it—that will entail the rationality of all and only the right epistemic rules, Modus Ponens and company.

Goldman, Boghossian, Enoch and Schechter, and Horwich don't offer explicit details about when a property is *deep* or *substantive* in the right way, or one *in virtue of which* another property holds, or one that figures in an *ur-norm*. Enoch and Schechter are even up front about the know-it-when-you-see-it aspect of a good explanation: as they say, a good theory will be 'one that we would be happy to accept as relevant to justification, one that presents the relevant methods in a rationally positive light'. None of this is grounds for criticizing the proposed theories; you don't need to work out a theory of theory-giving before you're allowed to give a theory of rationality. And, I agree with these theorists that we in any case do have a decent, intuitive grip on when a property is capable of doing the explanatory work a theorist of rationality is seeking.

It will be useful to illustrate this unificationist model with a quick summary of the contemporary, mainstream proposals for explaining the rationality of our deductive epistemic rules. All these proposed theories aim to identify some normatively deep feature possessed by just these rules, Modus Ponens and company, not Peirce's Rule and the rest. Peacocke (1993, 2004) and Bealer (2000) have said something roughly like 'These are the rules that are guaranteed to be reliable by virtue of the way we grasp the concepts they involve'. Peacocke calls this a metasemantic theory. In a relatedly metasemantic approach, Boghossian (2003) and Wedgwood (2011) have said something like 'These are the rules which we cannot antecedently question the reliability of, because accepting them is a precondition on our grasp of certain basic concepts.' Enoch and Schechter (2008) and Wright (2004) follow a pragmatic approach, one inspired by Reichenbach (1938); they say, 'We have a special epistemic entitlement (or maybe even obligation) to deliberate about the world, and these rules are the ones that offer us our only hope of fulfilling our deliberative entitlement (obligation)'. In general, these proposals say that all and only the rational epistemic rules have some special property F (where F = being guaranteed to be reliable in a certain way, or being impossible for us to antecedently question the reliability of, or discharging certain obligations we have), and identifying this feature F is supposed

to be explanatorily illuminating. It doesn't matter if *F* is explicitly normative; the general form shared by these explanations is not that of a reduction from the normative to the non-normative. These proposed theories purport to explain why Modus Ponens and certain other epistemic rules are rational via unification, via identifying some unifying property that just those epistemic rules have, even if it is a feature characterized in explicitly normative terms.

These theories all attempt to unify the epistemological data in just the same way that well-known moral theories attempt to unify the moral data. Think again of utilitarianism, the categorical imperative, ideal observer theories, or contractualist theories. And as mentioned, like in the moral case, all our best theories of rationality are known to face serious objections, and fail to enjoy anything remotely close to widespread acceptance.

Let me be clear that I *do* view unification as explanatorily illuminating. If some elegant, normatively deep, normatively fundamental feature *F* belonging to just the rational rules *were to be identified*, that *would* be extremely philosophically interesting; it *would* indeed be explanatorily illuminating. However, every proposed theory faces serious known objections, and none enjoys remotely widespread acceptance. So, the true theory of rationality, if it exists, is as yet undiscovered. And in the next section I'll offer my positive argument that it doesn't exist at all.

Horwich (2008) favors the view that there is no theory of rationality. He says: '[T]he most fundamental fact of epistemic normativity might very well be that a certain longish and diverse list of belief-forming procedures itemizes the rational ones. It would be entirely unsurprising, in other words, if the correctness of what we ordinarily regard as basic epistemic norms simply could not be explained' (p.471). (Note that this view Horwich and I favor does not go as far as particularism, the view of Dancy (2004). Particularism denies that we can even provide a 'longish' list that would itemize the procedures or rules of rational belief-formation.) Horwich's case for this view relies largely on argument by elimination: most of his paper is devoted to considering the mainstream proposed theories, and objecting to each one. This leads him to then suggest that epistemologists may have, perhaps by overgeneralizing from empirical science, too hastily presupposed that there is a theory of rationality out there to be discovered. I want to show a way to go beyond Horwich's negative argument and burden-shifting strategy.

3. No Property Unifies the Rational Rules

The previous section showed that a theory of rationality is an explanation of why certain epistemic rules are rational, an explanation that proceeds by identifying a unifying property of the rational rules. A theory of rational deduction would identify a property that unifies the rational epistemic rules for inferring deductive consequences of one's premises. In this final section of the paper, I'll argue that if communism is true, then there is no such theory of rational deduction.

This section has three sub-sections. Sub-section 3.1 looks at how deflationists about truth argued, on the basis of a premise about the function of the predicate 'true', that no property unifies the set of truths. Deflationists thus provide us with

a rare and valuable precedent: they provide a model for how to argue from the function of a word to a conclusion about the property expressed by that word, in particular to the conclusion that nothing unifies instances of that property. The model can be applied using the communist premise for 'rational' and the conclusion that no property unifies the rational rules. However, the deflationist's model makes for a weaker argument than we need to settle for. Sub-section 3.2 describes a way that the meaning of 'rational' is conventional (a way that the meaning of 'true' is not); and this conventionality, as I show in sub-section 3.3, allows us to strengthen the argument that nothing unifies the rational epistemic rules.

3.1 Comparison with Deflationism about Truth

Quine (1970) pointed out that the truth predicate serves a special function. Having the truth predicate allows us to make assertions that affirm (or deny) propositions that we are otherwise unable to affirm (deny). Specifically, it allows us to affirm propositions we can name or describe but cannot assert. A standard example: a layman might not be able to assert Fermat's Last Theorem. (In other words, the layman cannot assert the proposition that the equation $a^n + b^n = c^n$ has no solutions when a, b, c and n are natural numbers, and $n > 2$.) But, he can still affirm it by asserting the proposition that Fermat's Last Theorem is true. Another example: I must resort to using the truth predicate to affirm the infinitely many propositions that are instances of the law of non-contradiction.

Quine himself raised his point about the truth predicate's function in the context of arguing that propositions are dispensable, a longtime pet project of his. The far more interesting application of the point was made by several philosophers since Quine, most notably Leeds (1978), Horwich (1990/98), and Field (1994). These philosophers said that Quine's point, which as I've emphasized is about a word (or concept), shows something about the associated *property*, namely that there is no theory of it. Their reasoning was as follows.

They observed that the truth predicate is able to perform the function Quine identified purely because of a fact about the way we *use* the word. This use is just the following: we treat any attribution of truth to a proposition (or sentence) as cognitively equivalent to that proposition (sentence) itself. (Another way of describing the use that suffices to allow the truth predicate to perform its Quinean function is this: we accept every instance of the equivalence schema, namely *p is true if and only if p*. Exactly how to put it may depend on how you prefer to address the semantic paradoxes. For example, you might restrict our endorsement of the biconditionals of the schema, while still allowing inferences from either side of any admitted biconditional to the other.) Why do the deflationists think an observation about our use of a word reveals something about the associated property? They say: because, we can perfectly successfully use the truth predicate in that way, and thus it can perfectly successfully perform its sole function, *whether or not there exists a theory of the property of truth*. On this basis, they conclude that there is no theory of truth.[5]

The communist can emulate the deflationist's argument. 'Rational' can serve its coordinating function equally well regardless of whether there is a property unifying

all and only the rational rules. (For 'rational' to serve its function, the rational (deductive) rules need to be reliable (sound) and powerful (complete), but, as I've been emphasizing with the example of Peirce's Rule, this is not also a sufficient condition, and thus not a potential unifying condition, for the rationality of our rules.) So, to the extent that the deflationists about truth give a good argument, we have a good argument that nothing unifies the rational rules. If there is some necessary and sufficient property of the rational rules that unifies those rules, it is not because the function of 'rational' requires there to be such a property.

However, I *don't* want to offer this as the best argument that nothing unifies the rational rules. Although many are tempted by the deflationist's argument to accept its conclusion, I think the argument is not strong enough to rationally compel anyone who accepts its premises to accept its conclusion. I think this because I think it is not absurd to hold the combination of views that says both (i) the deflationist has the correct (and complete) view of the function of the predicate 'true', namely that the predicate serves a generalizing role, and (ii) the predicate's extension, the set of truths, is unified by some property they and they alone share.

In the next two sub-sections, I'll argue that it is, however, absurd to hold the analogous combination of views that both (i) the communist has the correct (and complete) view of the function of 'rational', and (ii) there is a property that unifies the rational rules. Sub-section 3.2 sets the argument up by elaborating a feature 'rational' has which 'true' does not have. Sub-section 3.3 then says why it is so absurd to think that 'rational' has such a feature yet the rational rules are still unifiable.

3.2 What We Call Rational is a Matter of Convention

Deflationists say that 'true' serves its generalizing function equally well whether or not a property unifies the things in its extension, i.e. the true things, i.e. the truths. Communists say that 'rational' likewise serves its coordinating function equally well whether or not a property unifies the things in its extension, i.e. the rational things, i.e. the rational epistemic rules. But, the function communists claim for 'rational' is also, in a further respect, indifferent to the nature of the things in its extension. Unlike the truth predicate, 'rational' can serve the communist's coordinative function equally well *as long as it is assigned any one of many possible extensions*.

The coordinative function is served when all members of an epistemic community share the same stock of epistemic rules, and those rules make up a reliable and powerful stock. Focusing just on our rules for deductive reasoning, what we want is a stock of rules that is sound and complete. In our actual community, we achieve this by coordinating upon Modus Ponens, Conditional Proof, Reductio, and some others; these are the epistemic rules that, very roughly, correspond to rules for derivations in a standard natural deduction proof theory. We promote these rules by assigning 'rational' an extension that includes just these rules for deductive reasoning. But, there are other stocks of rules for deductive reasoning that we could have coordinated upon, and that would have served our purposes equally well. To promote coordination upon such an alternative stock, we would assign

'rational' (or whatever our term for epistemic evaluation would be) an extension that contained just those rules for deductive reasoning.

The existence of alternative yet equivalent proof theories in logic indicates what such an alternative stock of rules can look like. In a number of elegant, sound and complete axiomatic formal systems, Peirce's Law is taken as an axiom. Tarski and Łukasiewicz give an example in Tarski (1983, chapter 4). This suggests how, as an alternative to our actual epistemic practice, a group could fruitfully engage in a practice that promoted following what I'm calling 'Peirce's Rule' as a rule of deductive reasoning. The promotion of Peirce's Rule would then involve including it in the extension of 'rational'. (As Harman (1986) points out, a rule of derivation in a proof theory is not the same thing as an epistemic rule, a rule for reasoning. I'm only saying the derivation rules in certain proof theories can serve as a suggestive model for what certain epistemic rules would look like.)

I'm claiming that there are alternative extensions for 'rational' that equally well serve the coordinative function of epistemic evaluation. I do not mean these alternatives are *under all circumstances* equally good for this. If we take it as given that our cognitive psychology is as it *actually* is, then of course different extensions for 'rational' would not serve the coordinative function equally well. Modus Ponens *is* psychologically privileged over Peirce's Rule. I am claiming that if we just consider the intrinsic features of the epistemic rules themselves, then Modus Ponens is not better than Peirce's Rule at serving our fundamental epistemic goal, our goal of true belief. With regard just to our epistemic goal, the choice between an extension for 'rational' that excludes Peirce's Rule versus an extension that includes it is thus an *arbitrary* choice.

A choice among arbitrary options results in a *convention*. The choice of what to call rational, the choice to include Modus Ponens but not Peirce's Rule in the extension of 'rational', thus results in a convention. But, I need to qualify my claim that *the choice* results in a convention. The choice is not *ours*. We cannot choose our cognitive psychological hardwiring, and given our hardwiring we are not then left with multiple equally good options to choose among for what to call rational. The choice of what epistemic rules to promote is thus not up to us in the sense that other conventional choices, such as which side of the road to drive on, are up to us. The arbitrary choice of which epistemic rules to promote is not a choice that's up to us, but is rather up to nature. But, I think that this is still a kind of convention. In defense of calling this a convention, I offer an analogy: think of the signaling convention used by honeybees, the waggle-dance. It's not up to the bees to use this or another equally good signaling system, but there is still an arbitrary choice here that nature made for the bees, and it strikes me as one we can appropriately call a kind of convention. In the same sense, I say it is a convention that we call Modus Ponens, but not Peirce's Rule, rational.[6]

3.3 The Candidate Extensions of 'Rational' aren't All Unifiable

I claim it is absurd, an untenable combination, to say both (i) that we call Modus Ponens, and not Peirce's Rule, rational as a matter of convention, in the sense just elaborated, and (ii) that some property unifies the rules that we actually call

rational, in the way a theory of rationality would have it. In this final sub-section, I'll argue for this claim.

Consider each of the many candidate extensions 'rational' could have that would equally well serve our epistemic goal of believing the truth. As noted, some of these candidates will include Peirce's Rule, though the actual extension doesn't. And now consider this question: do *all*, or even *most*, or even *many* of these candidates, these lists of epistemic rules for deductive reasoning, have their own unifying property? Is there, for each list of rules, a property that unifies all and only the rules on that list? I am asking about a property that would unify a list of epistemic rules, unify them in the way a theory of rationality—if there really were one—would identify a property that unifies the (actual) rational rules.

It is crucial here to emphasize, once again, that a *unifying* property of this sort, the sort sought by a theorist of rationality, is no mere necessary and sufficient property for appearing on the given list. It is not even enough to be a necessary and sufficient property that is also "fundamental", or "interesting", or "deep" in any old way. The property must be deep *in the right way*, in a way that makes it suitable for explaining the normative facts a theory of rationality is meant to explain. Recall Enoch and Schechter's warning: the right explanation needs to be 'one that we would be happy to accept as relevant to justification, one that presents the relevant methods in a rationally positive light.'

The various candidate extensions of 'rational', the various lists, after all do have necessary and sufficient properties that are fundamental, interesting, and deep in the wrong ways, ways that are irrelevant to the normative status of any epistemic rules. Logicians study such properties when they study the proof theories associated with these candidate lists of epistemic rules. For example, Łukasiewicz (1948–1950) found a sound and complete proof theory with the interesting feature of being the shortest, single-axiom proof theory. There are also lots of interesting properties that a list of rules actually, though not necessarily, possesses: this includes the psychologically, biologically, and historically significant properties that the actual rational rules have for us.

But, as was elaborated in section 2, a theory of rationality is supposed to identify a very special kind of necessary and sufficient property if it is to unify the rational rules, to unify them in the explanatorily illuminating way that is intended. The property must be not only fundamental, deep, or interesting, but *normatively* fundamental, deep, or interesting. I don't have a definition for when a property is normatively fundamental, deep, or interesting, but I do know lots of properties that don't fit the bill. Consisting of a single, very short, axiom is, on its own, decidedly not normatively fundamental. Nor, of course, is any merely psychologically significant property. Normative fundamentality is something else. The various proposed theories of rationality constitute rough stabs at the kind of normative fundamentality that theorists of rationality have been seeking. Recall section 2's summary of contemporary views. Peacocke and Bealer proposed the property of being a rule that is guaranteed to be reliable by the theory of concept possession. Boghossian and Wedgewood proposed the property of being a rule whose reliability cannot be antecedently questioned due to certain limits on concept possession. And Enoch

and Schechter and Wright proposed the property of fulfilling certain fundamental epistemic entitlements or obligations to deliberate about and explain the world. (Again, as the last example illustrates, it might even be that we have to use normative language to characterize the sought after normatively fundamental property.) I'm not claiming all these examples really are normatively fundamental; maybe they are not, but they are at least in the right ballpark, unlike, say, axiom length.

A property that is normatively fundamental, and thus at least *capable* of unifying whatever it is necessary and sufficient for in the intended sense of unification, is a rare sort of property. And this is where trouble is created by the fact that 'rational' has so many, perhaps infinitely many, candidate extensions (these candidates, once again, being illustrated by the various proof theories that a community could usefully model their deductive reasoning after). It is not plausible that every, most, or even many of these candidate extensions possess a *normatively* fundamental, deep or interesting necessary and sufficient property. It is thus not plausible that most or many of the candidate extensions are unifiable, in the sought after sense.

I am offering it as a plausible claim that there are not normatively fundamental properties out there unifying the various candidate extensions of 'rational'. I don't have an argument for this on the basis of other assumptions; this is an "unargued" premise for me. But I think it is plausible, at least when it is clearly understood, and I think it is certainly more plausible than an alternative assumption that the theorist of rationality takes for his own unargued premise.

Boghossian identifies the theorist's unargued premise as follows: '[I]t is hard to believe that the property of being warrant-transferring is simply a primitive property that an inference pattern either has or fails to have. Surely, if an inference pattern is warrant-transferring there must be some property by virtue of which it is warrant-transferring.' Boghossian (2003, p.239). (Also see the earlier quote of Enoch and Schechter, p.14 above.) My claim is that it is even harder to believe that, for each of many lists of candidate rules for deductive reasoning, there is something in virtue of which those rules appear on their respective lists; what I claim is more surely true is that there isn't such an appropriately fundamental property unifying each of many candidate extensions, that is, a *normatively* fundamental property. Of course, Boghossian, Enoch and Schechter and the other theorists were not thinking about epistemic communism or about multiple candidate extensions for 'rational'. But, once we take up the communist's view that it is a matter of convention which one of many candidate lists of rules we call rational, we have to ask how plausible it is, or how hard to believe it is, that many lists of epistemic rules are each unifiable.

And now, to complete the argument that there is no theory of rationality, remember that our actual extension for 'rational' is an arbitrarily selected one. Even supposing it should happen to be that there is a normatively fundamental property out there unifying some list of epistemic rules, we did not—or rather, nature did not—seek out such a list. The selection of our rules among the many candidates was arbitrary. And that means, even supposing there is a unifiable list of rules out there (which we don't even know), it would be a lucky fluke if our extension, the list of rational rules, hit that bull's-eye. It would be a fluke if *our* selected list, Modus Ponens and company, were unifiable by a normatively fundamental property.

My argument does not *prove* that nothing unifies our list of rational rules. I've only argued for the *low probability* of that. For all I've said, perhaps there is some argument out there that someone will come up with to show that one particular list of rules is thoroughly unified by some normatively fundamental property, and perhaps there is also at the same time an argument out there showing that the mechanics of evolution for psychology make it likely that intelligent creatures adopt that particular list of unifiable rules among the many candidate lists that can serve the coordinative function. What I've claimed is only that the chances of this are low right now, that is, given what we know right now. What we know right now is that it's quite hard to come up with any unifying properties for a list of rules (as demonstrated by epistemologists' struggles to come up with proposals for what might unify our list), and what we *don't* know yet is if the mechanics of evolution make likelier the adoption of any particular list of rules by creatures like us, and if so, why.

What should we do now? Where does all this leave the explanatory ambitions of epistemology? Although I am pessimistic about unifying our epistemic rules, I am very optimistic that we can still do interesting, and explanatorily illuminating, epistemology. I hope that epistemic communism, whether true or not, is itself a demonstration of this possibility. But, aside from a few outliers of the epistemological mainstream (Craig (1990) and Reynolds (2002) deserve mention), communism does not fit the standard form of an epistemological theory or explanation. I do not by any means wish to suggest communism exemplifies a new model that epistemological theories or explanations should now strictly conform to. My only suggestion is that we look beyond the traditional search for necessary and sufficient, much less unifying, properties of the rational epistemic rules, and, one way or another, continue to shake things up in epistemology.

Summary

Epistemic rules like Peirce's Rule raise a puzzle: why are some valid rules not called rational? Epistemic communism provides an answer: we use epistemic evaluations to promote coordination, which is valuable because it helps us to pursue true belief as a team, a team of parallel epistemic processors who can safely share their results through trustworthy testimony. This is the whole function of epistemic evaluations; in particular, there is no distinctive, first-personal function for them to play.

Extrapolation from recently proposed theories of rationality suggests that a theory would provide a property that is necessary and sufficient for being a rational rule, and would also unify the rational rules in an explanatorily illuminating way. Deflationists about truth said there is no reason to think there is some property that unifies the truths. I offered an even stronger argument that no property unifies the rational rules. There is no reason to believe that, for each, or even many, of the countless candidate extensions that would have allowed 'rational' to serve its coordinative function, there is a property unifying the rules in that extension. Since our extension was selected arbitrarily from among the candidates, there is no reason to think our rules have a unifying property.[7]

Notes

[1] See Goldman (1979/92, 1986) on reliablism. See Harman (1986); Huemer (2001, 2007) on conservatism. See Peacocke (1993, 2004); Bealer (2000); Boghossian (2003); Wedgwood (2011) on the metasemantic account. See Enoch and Schechter (2008); Wright (2004) on the pragmatic account. (Pp.4 and 14 provide a quick sketch of what these theories say.)

[2] See Hawthorne (2004); Stanley (2005); Fantl and McGrath (2009) for defenses of views that posit subject-sensitivity, though the forms of sensitivity discussed there do not concern coordination with the reasoner's community.

[3] See Cohen (1987); DeRose (2009); Lewis (1996) for contextualist views, though, again, the views discussed don't concern sensitivity to anything like coordination with one's community. Cohen (1987) has points of similarity to communism that most contextualist views don't have, in particular the emphasis on social groups, and even those groups' reasoning powers or abilities; see pp.15–16.

[4] My point here is reminiscent of an observation emphasized in Moran (2001), Shah (2003) and Shah and Velleman (2005). They observe that, whenever a thinker asks herself the deliberative question of whether to believe *p*, the answer to that question is settled by, and can only be settled by, the matter of whether *p*. What they say about the question of whether *to believe p* does not carry over perfectly to the question of whether *it's rational to believe p*. Insofar as epistemic akrasia is possible, one can answer the question of whether it's rational to believe *p* without answering the question of whether *p*. Still, under normal conditions, it seems that what they say about the question of whether-to-believe does carry over to our question of whether-it's-rational-to-believe.

[5] To back up this interpretation of the deflationist's argument, see section III of Leeds (1978) and Horwich (1990/98, p.5)

[6] Sider (2011, section 4.3) develops a notion of convention, which he calls *candidate-selection conventionality*, for assigning one among alternative meanings to a word that equally well serve the same function. Sider then offers, in section 4.4, a modification resulting in a notion he calls *subjectivity*; the modification eliminates the element of explicit choice that convention is traditionally thought to require. Sider's notion of subjectivity nicely captures the communist's view of 'rational'. (I think it's even open to dispute whether the classic accounts of convention found in Hume (1739–40/1978, p.490) and Lewis (1969, 1975) require that an agent have made a choice.)

References

Adler, Jonathan (2002). *Belief's Own Ethics*. Cambridge, MA: MIT Press.

Alston, William (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Bealer, George (2000). "A Theory of the A Priori." *Pacific Philosophical Quarterly* 81: 1–30.

Boghossian, Paul (2003). "Blind Reasoning." *Proceedings of the Aristotelian Society, Supplementary Volume* 77(1): 225–248.

BonJour, Laurence (1980). "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5: 135–150.

——— (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Burge, Tyler (1993). "Content Preservation." *The Philosophical Review* 102(4): 457–488.

Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Oxford University Press.

Cohen, Stewart (1987). "Knowledge, Context, and Social Standards." *Synthese* 73(1): 3–26.

Craig, Edward (1990). *Knowledge and the State of Nature*. Oxford: Oxford University Press.

Dancy, Jonathan (2004). *Ethics without Principles*. Oxford: Oxford University Press.

Darwin, Charles (1871/74). *The Descent of Man*. 2nd edition. Project Gutenberg Etext, www.gutenberg.org.

DeRose, Keith (2009). *The Case for Contextualism*. Oxford: Oxford University Press.

Dogramaci, Sinan (2012). "Reverse Engineering Epistemic Rationality." *Philosophy and Phenomenological Research* 84(3): 513–530.

Enoch, David and Joshua Schechter (2008). "How Are Basic Belief-Forming Methods Justified?" *Philosophy and Phenomenological Research* 76(3): 547–579.

Fantl, Jeremy and Matthew McGrath, (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Field, Hartry (1994). "Deflationist Views of Meaning and Content." *Mind* 103(411): 249–285.

Foley, Richard (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.

——— (1993). *Working Without a Net*. Oxford: Oxford University Press.

Fricker, Elizabeth (1994). "Against Gullibility." In B. K. Matilal and A. Chakrabarti (eds.), *Knowing From Words*, 125–61. Dordrecht: Kluwer Academic Publishers.

Friedman, Michael (1974). "Explanation and Scientific Understanding." *The Journal of Philosophy* 71(1): 5–19.

Gibbard, Allan (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.

Goldman, Alvin (1979/92). "What is Justified Belief?" In *Liasons: Philosophy Meets the Cognitive and Social Sciences*, 105–126. Cambridge, MA: MIT Press.

——— (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Harman, Gilbert (1986). *Change in View*. Cambridge, MA: MIT Press.

Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.

Horwich, Paul (1990/98). *Truth*. 2nd edition. Oxford: Oxford University Press.

——— (2008). "Ungrounded Reason." *The Journal of Philosophy* 105(9): 453–471.

Huemer, Michael (2001). *Skepticism and the Veil of Perception*. Lanham, Maryland: Rowman & Littlefield.

——— (2007). "Compassionate Phenomenal Conservatism." *Philosophy and Phenomenological Research* 74: 30–55.

Hume, David (1739–40/1978). *A Treatise of Human Nature*. Ed. L. A. Selby-Bigge and P. H. Nidditch. 2nd edition. Oxford: Oxford University Press.

——— (1748/2000). *An Enquiry Concerning Human Understanding*. Ed. Tom L. Beauchamp. Oxford: Oxford University Press.

Kitcher, Philip (1981). "Explanatory Unification." *Philosophy of Science* 48(4): 507–531.

——— (1989). "Explanatory Unification and the Causal Structure of the World." In Philip Kitcher and Wesley Salmon (eds.), *Minnesota Studies in the Philosophy of Science*, volume XIII, 410–505. Minneapolis: University of Minnesota Press.

Lackey, Jennifer (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

Leeds, Stephen (1978). "Theories of Truth and Reference." *Erkenntnis* 13(1): 111–129.

Lewis, David (1969). *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.

——— (1975). "Languages and Language." In Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, volume VII, 3–35. Minneapolis: University of Minnesota Press.

——— (1996). "Elusive Knowledge." *Australation Journal of Philosophy* 74(4): 549–567.

Łukasiewicz, Jan (1948–1950). "The Shortest Axiom of the Implicational Calculus of Propositions." *Proceedings of the Royal Irish Academy. Section A: Mathematical and Physical Sciences* 52: 25–33.

McLarty, Colin (2010). "What Does It Take to Prove Fermat's Last Theorem? Grothendieck and the Logic of Number Theory." *The Bulletin of Symbolic Logic* 16(3): 359–77.

Moran, Richard (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton, NJ: Princeton University Press.

Peacocke, Christopher (1993). "How Are A Priori Truths Possible?" *European Journal of Philosophy* 1: 175–199.

——— (2004). *The Realm of Reason*. Oxford: Oxford University Press.

Pollock, John and Joseph Cruz (1999). *Contemporary Theories of Knowledge*. Lanham, Maryland: Rowman & Littlefield.

Pronin, Emily, Daniel Lin and Lee Ross (2002). "The Bias Blind Spot: Perceptions of Bias in Self Versus Others." *Personality and Social Psychology Bulletin* 28(3): 369–381.

Quine, W. V. (1970). *Philosophy of Logic*. Cambridge, MA: Harvard University Press.

Reichenbach, Hans (1938). *Experience and Prediction*. Chicago: Chicago University Press.

Reid, Thomas (1780/1981). *An Inquiry into the Human Mind on the Principles of Common Sense*. Ed. Derek R. Brookes. University Park: Pennsylvania State University Press.

Reynolds, Steven (2002). "Testimony, Knowledge, and Epistemic Goals." *Philosophical Studies* 110(2): 139–161.

Shah, Nishi (2003). "How Truth Governs Belief." *The Philosophical Review* 112(4): 447–482.

Shah, Nishi and David Velleman (2005). "Doxastic Deliberation." *The Philosophical Review* 114(4): 497–534.

Sider, Theodore (2011). *Writing the Book of the World*. Oxford: Oxford University Press.

Sperber, Dan, Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgi and Deirdre Wilson (2010). "Epistemic Vigilence." *Mind and Language* 25(4): 359–393.

Stanley, Jason (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Tarski, Alfred (1983). *Logic, Semantics, Metamathematics: Papers from 1923 to 1938*. Ed. John Corcoran; trans. J. H. Woodger. 2nd edition. Indianapolis: Hackett.

Weatherson, Brian (manuscript). "Do Judgments Screen Evidence?"

Wedgwood, Ralph (2011). "Primitively Rational Belief-Forming Practices." In Reisner and Steglich-Petersen (eds.), *Reasons for Belief*, chapter 9. Cambridge: Cambridge University Press.

White, Roger (2006). "Problems for Dogmatism." *Philosophical Studies* 131: 525–557.

Wright, Crispin (2004). "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume* 78(1): 167–212.