Ideal Rationality and Logical Omniscience

Does rationality require logical omniscience? Our best formal theories of rationality imply that it does, but our ordinary evaluations of rationality seem to suggest otherwise. This paper aims to resolve the tension by arguing that our ordinary evaluations of rationality are not only consistent with the thesis that rationality requires logical omniscience, but also provide a compelling rationale for accepting this thesis in the first place.

This paper defends an account of apriori justification for beliefs about logic that is designed to explain the thesis that rationality requires logical omniscience. On this account, apriori justification for beliefs about logic has its source in *logical facts*, rather than *psychological facts* about experience, reasoning, or understanding. The proposal is that logical truths provide sources of apriori justification to believe those truths that is *indubitable*, *infallible*, and *indefeasible*. This account is designed to explain why rationality requires logical omniscience.

This account has important consequences for the epistemic role of experience in the logical domain. On this account, experience is not a source of apriori justification for beliefs about logic, although it can enable one to use apriori justification that one already has in forming justified beliefs about logic. Similarly, experience is not a source of evidence that defeats apriori justification for beliefs about logic, although it can disable one from using apriori justification in forming justified beliefs. In a slogan, the epistemic role of experience in the apriori domain is not a *justifying* role, but rather an *enabling and disabling* role.

As I will explain, this account has consequences for a wide range of current issues in epistemology, including the connection between justification and truth, the nature of epistemic idealization, the relationship between propositional and doxastic justification, the function of higher-order evidence of one's cognitive imperfection, and the possibility of rational dilemmas.

Here is the plan for the paper. Sections 1 & 2 explain and motivate the thesis that rationality requires logical omniscience. Section 3 proposes an account of

apriori justification for beliefs about logic that explains this rational requirement. Sections 4 & 5 respond to objections by drawing a distinction between ideal and non-ideal standards of rationality. Sections 6 & 7 develop an account of the enabling and disabling role of experience in the epistemology of logic. Sections 8 & 9 explore the implications of this account for the structure of ideal rationality. Section 10 concludes with a summary of the main claims of the paper.

1. Probabilistic Coherence

A defining feature of probabilistic theories of rationality is the requirement that one's credences, or degrees of confidence, must be *probabilistically coherent* in the sense that they conform to the axioms of the probability calculus:

(1) For every P, pr (P) ≥ 0.

(2) If P is a tautology [i.e. a necessary truth], then pr (P) = 1.

(3) If P and Q are mutually exclusive, then pr (P v Q) = pr (P) + pr (Q).[1]

The assumption is that the credences of a rational agent can be represented by means of a probability function that assigns a real number in the unit interval between 0 and 1 to every set of points in a probability space. An agent's credence in a proposition is given by the number assigned by the probability function to the set of points in the probability space in which the proposition is true. If one is certain that a proposition is true, then its probability is 1; if one is certain that it is false, then its probability is 0; and if one has some intermediate degree of confidence that it is true, then its probability is intermediate between 0 and 1.

Probabilistic coherence entails that one is certain of all necessary truths, since the probability of any necessary truth is 1. But what is the relevant notion of necessary truth? A necessary truth can be defined as one that is true at all points in

---

[1] This informal statement of the axioms is taken from Christensen (2004: 16).

probability space, but different interpretations of probability space yield different interpretations of probabilistic coherence.[2]

On a logical interpretation, the points in probability space are logical possibilities and necessary truths are logical truths. On this interpretation, probabilistic coherence requires that one is logically omniscient in the sense that one is certain of all logical truths, one is equally confident of all logically equivalent propositions, and one has at least as much confidence in each of the logical consequences of a proposition as one has in the proposition itself.[3] Arguably, however, the logical interpretation is too narrow, since the boundaries of logic are typically drawn in a way that excludes apriori truths of arithmetic, probability theory, epistemology, and so on. Why suppose that rationality requires being certain of logical truths but not these other apriori truths?

On a metaphysical interpretation, the points in probability space are metaphysical possibilities. On this interpretation, probabilistic coherence requires that one is certain of all metaphysically necessary truths. But while the logical interpretation seems too narrow, the metaphysical interpretation seems too broad. It is implausible that rationality requires one to be certain of Kripkean aposteriori metaphysical necessities: that Hesperus is Phosphorus, that water is H2O, and so on. Moreover, for any metaphysically contingent aposteriori truth that p, there is a metaphysically necessary aposteriori truth that it is actually the case that p. So, on this proposal, probabilistic coherence requires omniscience about the actual world.

David Chalmers (2011) proposes an epistemic interpretation on which the points in probability space are epistemic possibilities. An epistemic possibility is construed as a maximally specific proposition about the actual world that cannot be ruled out conclusively – that is, infallibly, indefeasibly, and indubitably – by an ideal process of apriori reasoning. On this interpretation, probabilistic coherence requires that one is certain of any truth that is epistemically necessary in the sense that it can

[2] See Hajek (MS) for discussion of the link between probability and modality in connection with the regularity thesis – that is, the thesis that if pr (P) = 1, then it is necessary that P.
[3] The requirement of logical omniscience is also a consequence of deductive theories of rationality that require logical consistency and closure.

be justified conclusively by ideal apriori reasoning. In a slogan, probabilistic coherence requires not just *logical omniscience*, but *apriori omniscience*.[4]

Apriori omniscience can be viewed as a generalization of logical omniscience. My own view is that the motivations for logical omniscience in section 2 extend to other apriori domains. However, it is beyond the scope of this paper to argue for the extension in any detail. So, for the purposes of this paper, I will restrict my focus to logical omniscience with the aim of extending the proposal in future work.[5]

Logical omniscience requirements are often regarded as the result of an *idealization* that is inherent in probabilistic theories of rationality. But what is the relevant kind of idealization? We can distinguish genuinely normative idealization from the non-normative idealization that is familiar in science. It can be convenient to ignore false predictions of a scientific theory when it sufficiently approximates towards the truth, especially when these are mere side effects of mathematical machinery that is otherwise useful or indispensable. To take just one well-known example, the Lotka-Volterra model in population ecology treats population abundance as continuous when in fact it is discrete.

Probabilistic theories of rationality are not scientific theories of human reasoning, but normative theories: they don't tell us how we *do* reason, but how we *ought* to reason.[6] Nevertheless, they are often thought to involve non-normative idealizations in much the way as scientific theories. On this view, logical omniscience requirements are idealizations in the sense that they make false claims about rationality that we can safely ignore because they are mere side effects of the mathematics of the probability calculus. The rival view defended in this paper is that logical omniscience requirements are idealizations in the normative sense that they make true claims about what ideal rationality consists in – that is, about how we

---

[4] Not all apriori truths are epistemically necessary in the sense that they are conclusively justified, including Hawthorne's (2002) examples of what he calls "deeply contingent apriori knowledge".

[5] I am particularly interested in the extension to apriori epistemic principles. See Titelbaum (forthcoming) for relevant discussion.

[6] See Kahneman, Slovic and Tversky (1982) for scientific evidence that we routinely violate probabilistic coherence.

ought to reason. I provide some initial motivation for this view in the next section before explaining and defending it further in later sections of the paper.

2. The Asymmetry Thesis

Rationality does not require omniscience across the board. The whole point of rational evaluation is to determine how well agents respond to the limited evidence in their possession. Rational agents respond appropriately to their evidence even when it is limited in ways that result in ignorance and error about the facts.

Given that rationality does not require omniscience across the board, those who endorse logical omniscience are committed to an asymmetry thesis:

> *The Asymmetry Thesis*: rationality requires omniscience and infallibility in logical domains but not empirical domains.

But what could motivate this asymmetry between logical and empirical domains? One prominent objection to probabilistic theories of rationality is that there is no principled motivation for the asymmetry thesis. Richard Foley gives a particularly clear statement of this objection:

> If a logically omniscience perspective…is an ideal perspective, one to which we aspire and one that we can do a better or worse job of approximating, so too is an empirically omniscient perspective. If this were a reason to regard all departures from logical omniscience as departures from ideal rationality, it would be an equally good reason to regard all departures from empirical omniscience as departures from ideal rationality. But of course no one wants to assert this. (Foley 1993: 161)

The aim of this section is to address Foley's objection by providing some intuitive and theoretical motivation for the asymmetry thesis.

The objection assumes that we should give a unified epistemology for logical and empirical domains. In fact, however, I want to argue that theoretical unification

is not a virtue in this instance, but a procrustean bed. There are structural differences between the epistemology of logical and empirical domains that are distorted by a unified theory.

I begin with an intuitive contrast. It is a familiar idea that a Cartesian demon could violate the connection between what seems true and what is true and thereby induce massive ignorance and error in the empirical domain without thereby impugning our rationality. But we cannot easily make sense of this in the logical domain. A Cartesian demon who induces massive ignorance and error by violating the connection between what is true and what seems true in the logical domain does not thereby leave our rationality intact.[7]

To begin with the case of error, consider the demon's victim who believes that contradictions are true, that affirming the consequent is valid, that conjunctions are false when their conjuncts are true, and so on. Moreover, she believes these things just because they seem true. Assuming symmetry between logical and empirical domains, we should say that these beliefs are fully rational and yet mistaken. But that just seems wrong. Intuitively, the demon has not deceived the subject in a way that leaves her rationality intact. Rather, by making these absurdities seem true, the demon has compromised her rationality. Rationality is inconsistent with too much logical error.

Similar considerations apply in the case of ignorance. Consider the demon's victim who is agnostic about whether contradictions are false, whether affirming the consequent is invalid, whether conjunctions are true when their conjuncts are true, and so on. She does not believe these things, even when she considers them, because they do not seem true. Once again, it just seems wrong to say that her logical ignorance leaves her rationality intact in the same way that sensory blindness might result in empirical ignorance about the environment without thereby compromising rationality. Rationality is inconsistent with too much logical ignorance.

---

[7] I am indebted here to discussion with Roger White. See also Jarvis and Ichikawa (2013: Ch. 12-13) for related discussion in the context of their critique of what they call 'experiential rationalism'.

So there is an asymmetry between logical and empirical domains. In the logical domain, but not the empirical domain, a Cartesian demon could induce massive ignorance and error by violating the connection between what seems true and what is true without thereby compromising one's rationality. Even so, the asymmetry thesis might seem like overkill insofar as it is committed to the thesis that rationality requires logical omniscience. After all, we don't always regard violations of logical omniscience as irrational.

In section 5, I'll respond to this objection by drawing a distinction between ordinary standards of rationality that take into account our human limitations and ideal standards of rationality that abstract away from them. Logical omniscience is required by the very demanding standards of ideal rationality, but not the more relaxed standards of ordinary rationality. Ordinary standards of rationality are consistent with some degree of logical ignorance and error even if there are limits on how far it can extend. But logical ignorance and error always constitutes some departure from ideal standards of rationality.

The intuitive motivation for the asymmetry thesis is not irresistible. One might follow Huemer (2001) in maintaining that the victim of deception can be fully rational and yet massively mistaken about logic. Or one might follow Williamson (2007) in denying that the victim of deception can be fully rational while massively mistaken about the empirical domain. But these responses threaten to obscure the intuitive differences between logical and empirical domains.

Even so, one might reasonably doubt how much weight these intuitive considerations can bear. In the absence of a more principled theoretical motivation for the asymmetry thesis, the relevant intuitions might be outweighed by the theoretical advantages of unifying the epistemology of logical and empirical domains. Arguably, however, these intuitive considerations do point in the direction of a more principled theoretical motivation for the asymmetry thesis.

Christensen (2004) proposes an argument for a version of the asymmetry thesis in the following passage:

Logical omniscience emerges naturally as the limiting case of one of the basic ingredients of good thinking, in a way that empirical omniscience does not. We know that certain structural aspects of the claims we believe have a bearing on their possible truths (e.g. a conjunction is true only if each of its conjuncts is true). Formal logic studies these relationships. It seems clear that many instances of bad thinking involve failing to respect these relationships (we should not believe a conjunction more strongly than its conjuncts). Eliminating this sort of mistake yields, in the limit, logical omniscience. Given that no such result holds for empirical omniscience, it seems to me that we have a clear motivation for treating the two differently when theorizing about rationality. (2004: 156-7)

Christensen illustrates the point with the following example. Kelly knows that she is near a grizzly bear cub and that anyone near a grizzly bear cub is in danger. But she doesn't "put two and two together" and so she fails to conclude that she is in danger. Meanwhile, Cherry also knows that anyone near a grizzly bear cub is in danger, but she doesn't know that Kelly is near a grizzly bear cub, and so she doesn't infer that Kelly is in danger. Intuitively, Kelly is rationally defective because she fails to respect logical relationships, but Cherry is not since she is merely ignorant of an empirical fact. Christensen concludes, "There is a clear intuitive basis in our ordinary conception of rationality for distinguishing logical lapses from ordinary cases of factual ignorance" (2004: 155).

Christensen's argument for the asymmetry thesis can be reconstructed as follows. The first step is that rationality is an ideal of good reasoning: that is to say, being ideally rational requires reasoning well. The second step is that reasoning well in the relevant sense does not require *empirical omniscience* in the sense that one is certain of all empirical truths. Good reasoning is perfectly consistent with empirical ignorance and error. The third step is that reasoning well in the relevant sense requires *logical omniscience* in the sense that one is certain of all logical truths. Good reasoning is not consistent in the same way with logical ignorance and

8

error. Hence, the conclusion follows that rationality requires logical but not empirical omniscience.

So the asymmetry thesis is crucial for preserving a rational ideal of good reasoning. Uncertainty and error about logic has negative ramifications for the quality of one's reasoning. For instance, if I doubt that modus ponens is valid, I might be very confident that I'm near a bear cub, while being much less confident that I'm in danger, despite being certain that if I'm near a bear cub, then I'm in danger. Similarly, if I'm confident that affirming the consequent is valid, then I might be very confident that I'm near a bear cub on the basis of my confidence that I'm in danger together with my certainty that if I'm near a bear cub, then I'm in danger. The general point is that good reasoning is contaminated by uncertainty and error about logic. Since rationality is a matter of reasoning well, it follows that ideal rationality requires omniscience and infallibility about logic.

A correlative point is that good reasoning leads rationally to true and certain beliefs about logic. For example, the validity of modus ponens can be ascertained by reasoning in accordance with modus ponens and the invalidity of affirming the consequent can be ascertained in much the same way.[8]

One might object that we need to distinguish more sharply between logical reasoning on the one hand and beliefs about logic on the other. For instance, we can imagine someone who reasons well without forming beliefs about logic at all. Alternatively, we can imagine someone who reasons well despite having false beliefs about logic because their reasoning does not take those beliefs into account. Arguably, however, reasoning well by ideal standards of rationality involves reasoning one's way towards beliefs about logic and then integrating one's logical reasoning with one's beliefs about logic.

I conclude that the asymmetry thesis is initially well motivated by a combination of intuitive and theoretical considerations. Still, the asymmetry thesis stands in need of further explanation. This is the topic of section 3.

---

[8] See Boghossian (2000) in defense of rule-circular justification for logical beliefs.

3. Explaining the Asymmetry Thesis

My aim in this section is to propose an account of apriori justification for beliefs about logic that is designed to explain why the asymmetry thesis is true. What explains why rationality requires omniscience and infallibility in logical domains but not empirical domains? My working assumption is that facts about rationality are explained by facts about reasons or justification. More specifically, I assume that rationality requires one to believe a proposition just when and because one has sufficient reason or justification to believe that proposition. And one complies with the requirement just when and because one believes the proposition on the basis of sufficient reasons or justification. If rationality requires one to be certain of all logical truths, then it follows that one has sufficient reason or justification to do so. But then what could be the source of this justification?

The explanation that I propose appeals to an asymmetry in the sources of reasons or justification in logical and empirical domains. Aposteriori justification for empirical beliefs has its source in experience, whereas apriori justification for logical beliefs has its source in the logical facts themselves. On this view, the logical domain has a fact-based epistemology, whereas the empirical domain has an experience-based epistemology. The asymmetry thesis therefore has important consequences for the role of experience in the epistemology of logic.

Apriori justification is traditionally defined as justification that does not have its source in experience.[9] But the term 'experience' can be understood in a broad sense or a narrow sense. The narrow sense of the term refers only to perceptual experience in the various sensory modalities. But some philosophers claim that there is cognitive experience associated with understanding, reasoning, and intuition that is sui generis and irreducible to perceptual experience. I use the term 'experience' in a broad sense that includes any mental state or episode that is conscious in the phenomenal sense that there is "something it is like" for the subject to undergo that mental state or episode. Experience in this sense is a genus whose species include both perceptual experience and cognitive experience.

---

[9] See Casullo (2003: Ch. 2) for the traditional definition and discussion of the distinction between broad and narrow conceptions of experience.

What is the epistemic role of experience so construed? More specifically, does experience play the same kind of epistemic role in apriori and aposteriori domains? In what follows, I will compare and contrast two opposing answers to this question. *Experientialism* is the view that all justification has its source in experience and hence the distinction between apriori and aposteriori justification is explained in terms of a distinction between experiential sources: aposteriori justification has its source in perceptual experience, whereas apriori justification has its source in cognitive experience. *Non-experientialism* is the view that the distinction between apriori and aposteriori justification is explained in terms of a distinction between experiential and non-experiential sources: aposteriori justification has its source in experience, whereas apriori justification does not. On this view, not all justification has its source in experience.

Experientialism can be illustrated with reference to Huemer's (2001: 99) principle of phenomenal conservatism:

> *Phenomenal Conservatism*: if one has an experience in which it seems to one that *p*, then one thereby has defeasible justification to believe that *p*.

One of the attractions of phenomenal conservatism is that it promises to unify the epistemic roles of perceptual experience and cognitive experience by subsuming them under a single epistemic principle. As stated, phenomenal conservatism states only a sufficient condition for defeasible justification, but it can be strengthened to yield conditions that are both necessary and sufficient for defeasible justification:

> *Strengthened Phenomenal Conservatism*: one has defeasible justification to believe that *p* iff one has an experience in which it seems to one that *p*.

The result is a version of experientialism on which all justification has its source in experiential seemings. According to Huemer, the category of experiential seemings includes not only sensory or perceptual seemings, but also cognitive or intellectual seemings. So the distinction between apriori and aposteriori justification can be

defined in terms of a distinction among experiential seemings: aposteriori justification is aposteriori has its source in perceptual experience, while apriori justification has its source in cognitive experience.[10]

Experientialism is very much in tension with the asymmetry thesis. According to experientialism, all justification has its source in experience. However, experience is not perfectly calibrated with the facts. There is no guarantee either in the case of perception or cognition that what seems to be true is in fact true or vice versa. Experientialism is therefore consistent with the rationality of widespread ignorance and error in apriori and aposteriori domains alike.

One might supplement experientialism with the further condition that an experiential seeming provides justification to believe a proposition only if it is true. Moreover, one might stipulate that the justification provided by these experiential seemings is infallible, indefeasible, and indubitable. But this sacrifices the unity between apriori and aposteriori domains and thereby undercuts much of the motivation for experientialism. Moreover, it fails to explain the asymmetry thesis, since logical truths need not seem true, so it preserves logical infallibility at the cost of abandoning logical omniscience.

To explain the asymmetry thesis, I propose a version of non-experientialism on which apriori justification for beliefs about logic has its source not in facts about experience, but rather in the logical facts that one has apriori justification to believe. On this view, logical truths themselves provide apriori justification that is infallible, indefeasible, and indubitable. Following Alston (1971), I call this a *truth-sufficiency* account of a priori logical justification.[11]

There is a foundationalist version of this proposal that draws a distinction between basic and non-basic logical truths, or axioms and theorems. On this version, a priori justification to believe the theorems has its source in a priori justification to

---

[10] See BonJour (1998), Huemer (2001), Chudnoff (2013), and Bengson (forthcoming) for versions of experientialism.

[11] Alston considers truth-sufficiency as an account of privileged introspective access to one's mental states. See Smithies (2012a) for a truth-sufficiency account of introspective justification that provides a model for the truth-sufficiency account of apriori logical justification proposed here.

believe the axioms, but one's a priori justification to believe the axioms has its source in the truth of those axioms. For current purposes, I'll remain neutral about whether or not this foundationalist version is correct.[12]

The asymmetry thesis is explained by an asymmetry in the sources of justification in logical and empirical domains. Aposteriori justification for empirical beliefs has its source in the way things seem, whereas apriori justification for logical beliefs has its source in the way things are. This explains why mismatches between what seems true and what is true constitute rational failings in the logical domain but not the empirical domain.

In the empirical domain, ideal rationality is a matter of calibrating one's beliefs with the experiences that are the sources of one's aposteriori justification for empirical beliefs. Since these experiences are not perfectly calibrated with the facts, ideal rationality is compatible with ignorance and error in the empirical domain. In the logical domain, by contrast, ideal rationality is a matter of calibrating one's beliefs with the logical facts that are the sources of apriori justification for logical beliefs. Since this apriori justification is infallible, indefeasible and indubitable, ideal rationality is therefore incompatible with ignorance and error in the logical domain.

In the remaining sections of this paper, I'll develop this proposal in response to a series of objections.

4. Rationality and Human Limitations

The first objection I want to consider begins from the observation that complying with the requirement of logical omniscience is beyond our limited human capacities. And yet, the objection continues, we cannot be bound by a requirement of rationality unless we are capable of complying with it. Therefore, it follows that logical omniscience is not a requirement of rationality after all.

Here is a more detailed formulation of the objection that combines a deontological conception of rationality with a rational 'ought' implies 'can' principle:

---

[12] This may have been Frege's view. See Burge (1998) for discussion.

(1) If rationality requires logical omniscience, then I ought to comply with this requirement. (The deontological conception of rationality.)

(2) If I ought to comply with the requirement of logical omniscience, then I am capable of complying with it. (The rational 'ought' implies 'can' principle.)

(3) Therefore, if rationality requires logical omniscience, then I am capable of complying with this requirement.

(4) I am not able to comply with the requirement of logical omniscience.

(5) Therefore, rationality does not require logical omniscience.

Premise (4) is hard to dispute. There are infinitely many logical truths and, since we are finite creatures, we cannot grasp all of them. One might respond that there are infinitely many truths such that a finite creature can grasp each of them individually even if they cannot grasp all of them collectively. But there are no finite limits on the complexity of logical truths and so there will be always be some that are beyond the grasp of any finite creature. One might restrict the scope of the logical omniscience requirement to propositions that the thinker can grasp.[13] But the problem remains because grasping a proposition is not sufficient for ascertaining whether it is true. It is not within normal human capacity to infallibly determine the logical truths. Moreover, the logical truths are not decidable, so there is in principle no algorithm that could generate all of them.

The interim conclusion (3) is much less plausible. It is vulnerable to the objection that avoiding epistemic irrationality is sometimes beyond the capacity of individual human subjects. Consider patients with Capgras delusion who believes their spouse has been replaced by an imposter or patients with Cotard delusion who believes they are dead. These delusional patients are incapable of complying with the requirements of irrationality.[14] One might reply that this is within normal

---

[13] We need some restriction because every concept can figure in some logical truth, but rationality does not require grasping every concept. So the scope of the logical omniscience requirement should be restricted to propositions composed from concepts that the subject actually grasps or would if she were ideally rational.

[14] See Feldman and Conee (1985: 17), Alston (1988: 286-8), Pryor (2001: 114-5), and Christensen (2004: 161-2) for additional examples of this kind.

human capacity even if it is beyond the capacity of delusional patients. But if there are requirements of rationality that apply to individual human beings who are incapable of meeting them, then why can't there be requirements of rationality that are beyond human capacities in general?

In general, there is no good reason to suppose that all evaluative ideals worth caring about must be humanly attainable. We can and do evaluate the performance of human beings along dimensions whose extremes lie beyond human reach. Ideals of rationality – like ideals of morality, scientific understanding, and chess – may lie beyond our limited human capacities. The ideals themselves need not be humanly achievable so long as we can make sense of better and worse approximation towards those ideals. Christensen (2004: 162) puts the point effectively: "Not all evaluation need be circumscribed by the abilities of the evaluated. In epistemology, as in various other areas, we need not grade on effort."

If the interim conclusion (3) is false, then so is premise (1) or premise (2). But which of these premises should we reject? Does rationality impose an epistemic obligation upon us to be logically omniscient and, if so, does it follow that we must be capable of discharging this epistemic obligation?

Following Alston (1988), we might reject premise (1) by abandoning the deontological conception of rationality as a source of requirements that we are obliged to comply with and replacing it with an evaluative conception of rationality as a source of evaluative ideals. On this view, there is no obligation to comply with the rational requirement of logical omniscience, but there is nevertheless a distinctive kind of epistemic value that is associated with compliance.

An alternative reaction is to accept the deontological conception of rationality as stated in premise (1) but to reject the 'ought' implies 'can' principle as stated in premise (2). Feldman (2000) argues that there are "role-oughts" that apply to anyone who occupies a certain role regardless of how well they are able to play that role – for instance, chefs ought to make delicious food and jugglers ought to keep their balls in the air regardless of whether they are capable of doing so. Similarly, there are epistemic 'oughts' that apply to us in virtue of our role as

15

believers. As Feldman (2000: 676) puts the point: "It is our plight to be believers. We ought to do it right. It doesn't matter that in some cases we are unable to do so."

Deciding between these options depends on how we understand the relationship between values and obligations. Are we obliged to achieve evaluative ideals or merely to approximate towards them as closely as we can? Perhaps we can recognize a "thin" sense in which we are obliged to achieve ideals regardless of whether we are capable of doing so. According to Feldman and Conee, for example, "In any case of a standard for conduct...it is appropriate to speak of 'requirements' or 'obligations' that the standard imposes" (1985: 19). But we can also recognize a "thick" sense in which we have obligations that we can be held responsible for discharging in a sense that makes it appropriate to adopt reactive attitudes, such as praise and blame. Our obligations in this thick sense depend upon the limitations of our capacities – that is to say, we are obliged only to approximate towards ideals insofar as we are capable of doing so.[15]

The argument equivocates between thick and thin senses of 'ought'. In the thin sense, premise (1) is true, but premise (2) is false, whereas in the thick sense, premise (1) is false, but premise (2) is true. But there is no reading of the argument on which both of its premises are true. However we interpret the argument, it fails to establish that the requirements of rationality must be humanly attainable.

5. Ideal Versus Non-Ideal Rationality

The thesis that rationality requires logical omniscience also conflicts with our ordinary evaluations of rationality. We don't always regard violations of logical omniscience as irrational. In fact, we often regard them as eminently reasonable. Christensen illustrates the point with the following example:

> Suppose I prove a somewhat complex theorem of logic. I've checked the proof several times, and I'm extremely confident about it. Still, it might seem quite reasonable for me to be somewhat less than 100% confident. I should

---

[15] Compare Pryor (2001: 115, fn. 36) for a related distinction between thick and thin senses of 'obligation'.

not, for example, bet my house against a nickel that the proof is correct. (Christensen 2007: 8)

If it is reasonable for me to be certain that the proof is correct, then it is reasonable for me to bet my house against a nickel in the certainty that I will win. Intuitively, however, it would be quite unreasonable for me to accept this bet because it is much too risky. I know that I'm prone towards mistakes and hence that adopting a policy of accepting such risky bets is liable in the long run to leave me homeless with nothing but a pocket full of small change. It is therefore unreasonable for me to be certain that the proof is correct.

My response, as indicated in section 2, is to distinguish between ordinary standards of rationality that take into account our contingent human limitations and ideal standards of rationality that abstract away from them.[16] In Christensen's example, ideal standards of rationality require accepting the bet, whereas ordinary standards require rejecting the bet. Since these requirements conflict, one might ask which we ought to follow. But this question has no context-dependent answer: it depends on whether or not one is taking our contingent human limitations into account. In the terminology of section 4, there is a "thin" sense in which one ought to comply with ideal standards of rationality, but there is also a "thick" sense in which one ought to comply with ordinary standards. Our ordinary evaluations reflect standards of rationality that we can be held responsible for meeting in a sense that makes it appropriate to adopt reactive attitudes such as praise and blame.

One might be suspicious that this distinction undermines much of the interest of the thesis that rationality requires logical omniscience. If our ordinary evaluations of rationality reflect non-ideal standards, then why should we care about ideal standards? They might be relevant for ideal agents, but it is not clear what relevance they have for non-ideal agents like us. My response is that our ordinary standards of rationality are best understood as approximations towards

---

[16] See Schoenfield (2012) for a related distinction between two notions of rationality: what your evidence supports versus what you ought to believe given your evidence. See also Schoenfield (MS) for further development of her proposal.

17

ideal standards that take our contingent human limitations into account. But it is a further question how this approximation is to be understood.

One model says that one's credences approximate more closely towards ideal rationality insofar as they are closer to the credences that ideal rationality requires. But this model is clearly problematic. Suppose you give me a list of 50 theorems and 50 anti-theorems that are so complicated that I'm at chance when it comes to picking between them. In this forced-choice scenario, it is not reasonable by ordinary standards to be more confident of the theorems and less confident of the anti-theorems. Instead, it is reasonable to remain have a credence of 0.5 in each one. So the approximation has to be understood in a different way.

A better model says that one's credences approximate more closely towards ideal rationality insofar as they are held on the basis of better reasoning – that is, reasoning that approximates more closely towards ideal reasoning.[17] In the forced-choiced scenario, it is not reasonable to be more confident of the theorems and less confident of the anti-theorems because this would manifest bad reasoning. Since you're at chance when it comes to picking between them, reasoning in this way is no better than blind guesswork. Even if by some coincidence your credences align closely with the requirements of ideal rationality, they are not close to ideally rational because of the bad reasoning on the basis of which they are held.

If options are ranked by the standards of ideal rationality, then ideal rationality requires taking the options that are optimal in the ranking. But conforming to the requirements of ideal rationality is sometimes beyond our limited human capacities. The options that are optimal in the ideal ranking may be beyond our capacities in the sense that we are unable to take these options on the basis of good enough reasoning. Ordinary standards of rationality do not require us to take options that are beyond our capacities in this sense. The options required by ordinary standards of rationality are the best options that remain in the ideal ranking when those that are beyond our capacities have been excluded.[18]

---

[17] See section 6 for an account of good reasoning.
[18] On Kratzer's (1981) semantics, 'S ought to do A,' is true in a context C if S does A in all of the possible worlds in the contextually relevant modal base that are optimal

To illustrate the point with an example from the practical domain, consider Gary Watson's (1975) example of the squash player who is defeated in humiliating fashion. The most reasonable option by ideal standards would be to shake the opponent's hand and to thank him for the game. But this is not an available option for the squash player since he is so aggravated that he is liable to lash out with his racquet in anger. Even if he manages to shake hands without lashing out, it would seem unreasonable for him to do so, since he is in such danger of losing his temper. Given the human limitations that are salient in the context, the most reasonable option by ordinary standards is to leave the court immediately.

Similarly, in Christensen's example, the most reasonable option by ideal standards is to accept the bet in the certain knowledge that one will win a nickel. But this is not an available option for non-ideal agents like you and me. We might accept the bet and win, but we cannot do so on the basis of good enough reasoning to make this anything less than an unreasonably risky gamble. Given the human limitations that are salient in the context, the most reasonable option by ordinary standards is to reject the bet.

6. The Enabling Role of Experience

The version of non-experientialism that I have proposed faces a challenge to explain the role of experience in the epistemology of logic. Intuitively, someone who has an experience of proving a logical theorem or intuiting the truth of a logical axiom is thereby in a better epistemic position to acquire justified belief than someone who merely guesses that it is true. But if experience is not a source of apriori justification, then what explains its impact on the subject's epistemic position?

To illustrate the point, consider two subjects: Good and Bad. Good believes some logical theorem on the basis of an experience of proving that it is true,

---

according to the contextually relevant ordering source. Using Kratzer's framework, we can model ideal and non-ideal standards of rationality using the same ordering source but a different modal base. In contexts where ideal standards are relevant, the modal base includes worlds in which the subject takes options that are beyond her actual capacities, whereas these worlds are excluded in contexts were ordinary standards are relevant.

whereas Bad has no relevant experience of this kind and believes it on the basis of blind guesswork. There is surely some epistemic disparity between Good and Bad that stands in need of explanation. But if Good and Bad have equal justification to believe the logical theorem, then how can we explain the epistemic disparity between them?

I claim that the epistemic disparity between Good and Bad is best explained in terms of a distinction between good and bad reasoning. But the point can be made more precisely in terms of the distinction between *propositional justification* and *doxastic justification*. Propositional justification is a matter of having reasons or justification to believe a proposition, whereas doxastic justification is a matter of believing a proposition in a way that is reasonable, rational, or justified. Doxastic justification entails propositional justification, but not vice versa, since one can have justification to believe a proposition that one does not believe at all or that one does not believe in a way that is justified. A justified belief is one that is held on the basis of *good reasoning* – that is, it is held on the basis of doxastic dispositions that are sufficiently sensitive to the presence of propositional justification. Good reasoning – also sometimes called "proper basing" – is what converts propositional justification into doxastic justification.[19]

It is not controversial that there is a distinction between propositional and doxastic justification. A more controversial issue is whether there are doxastic constraints on propositional justification. Can a subject have propositional justification for a belief without having the capacity to convert it into doxastic justification by forming the belief on the basis of good reasoning? I agree with Feldman and Conee on this point:

Suppose that there were occasions when forming the attitude that best fits a

---

[19] See Kvanvig and Menzel (1990) for the proposal that doxastic justification is propositional justification plus basing. Turri (2010) proposes counterexamples, but as far as I can see, they do not satisfy the proper basing condition proposed here. But I am not sure that the proper basing condition can be understood independently of doxastic justification and so immunity from counterexample may be gained at the cost of reduction.

person's evidence was beyond normal cognitive limits. This would still be the attitude justified by the person's evidence. If the person had normal abilities, then he would be in the unfortunate position of being unable to do what is justified. (1985: 19)

A person's evidence may justify believing a proposition although he is unable to believe the proposition in a way that is appropriately based on the evidence.

To illustrate this point, consider the problem of the speckled hen.[20] If my experience represents that the hen has 48 speckles in the absence of defeaters, then I have propositional justification to believe that the hen has 48 speckles. Even so, if I were to believe that the hen has 48 speckles, then my belief would be doxastically unjustified. This is because my belief is held on the basis of doxastic dispositions that are not appropriately sensitive to the presence of propositional justification, since they tend to yield beliefs in the absence of propositional justification. For instance, I could easily have believed on the basis of the same doxastic dispositions that the hen has 47 or 49 speckles when my experience represents that it has 48 speckles, or that it has 48 speckles when my experience represents that it has 47 or 49 speckles. Therefore, my belief does not satisfy the proper basing condition for doxastic justification.

Now consider a case in which I believe some complicated logical theorem T on the basis of sheer guesswork. If T is true, then I have propositional justification to believe that T is true. Nevertheless, if I were to believe T on the basis of blind guesswork, then my belief would be doxastically unjustified. Once again, this is because my belief is held on the basis of doxastic dispositions that are not appropriately sensitive to the presence of propositional justification, since they tend to yield beliefs in the absence of propositional justification. For instance, I could easily have believed on the basis those very same doxastic dispositions some proposition T* that is not a logical truth, but a logical falsehood. Therefore, my belief does not satisfy the proper basing condition for doxastic justification.

[20] See Sosa (2003) and Smithies (2012b) for more detailed discussion of the problem of the speckled hen and its implications for the basing relation.

Compare Sosa's (2003) proposal that one knows that *p* only if one's belief is *safe from error* in the sense that one could not easily falsely believe that *p*. The safety condition is designed to explain why a belief cannot be knowledge if it just happens to be true. But it fails to explain why a true logical belief falls short of knowledge when it is held on the basis of blind guesswork. Such a belief is safe from error because its content is necessarily true, but it does not thereby meet the conditions for knowledge. Sosa's response is that such a belief is formed on the basis of doxastic dispositions that tend to yield false beliefs that are not safe from error. In Sosa's terminology, the belief is not *virtuous* in the sense that "one's belief derives from a way of forming beliefs that is an intellectual virtue, one that in our normal situation for forming such beliefs would tend strongly enough to give us beliefs that are safe" (2003: 290).

What is needed for doxastic justification is not safety from error, but rather safety from the absence of propositional justification. In the empirical domain, my beliefs can be doxastically justified even if they are false, and so not safe from error, so long as they are safe from the absence of propositional justification. In the logical domain, however, there is no distinction to be drawn between safety from error and safety from the absence of propositional justification, since apriori justification for beliefs about logic has its source in the logical facts themselves.

This means that true logical beliefs formed on the basis of blind guesswork are safe from error, and safe from the absence of propositional justification, since their contents are necessarily true. But they are formed on the basis of doxastic dispositions that also tend to yield beliefs in the absence of propositional justification. So these beliefs are not "virtuous" in the sense that Sosa defines. This is why they do not satisfy the proper basing condition for doxastic justification.

On this account, good reasoning or "proper basing" in the logical domain is really just a matter of reliability: it is a matter of forming beliefs on the basis of doxastic dispositions that are sufficiently sensitive to logical truth and safe from logical error. So, in principle, a reliable faculty of clairvoyance can yield justified beliefs about logic in the absence of any experience of intuition or proof. In practice, however, human reasoning about logic is much more reliable when it is based on

such experiences. So experience plays a crucial role in enabling human subjects to engage in good reasoning about logic.

Now we can explain the epistemic disparity between Good and Bad. Both Good and Bad believe the theorem, but only Good believes it on the basis of good enough reasoning. After all, Good believes it on the basis of an experience of proving the theorem, whereas Bad believes it on the basis of blind guesswork. So Good's belief is held on the basis of doxastic dispositions that are more reliably sensitive to the logical facts. Note that sensitivity comes in degrees. The claim is not that Good's experiences of proof or intuition are perfectly sensitive to the logical facts but merely that they are more sensitive than blind guesswork or wishful thinking. Therefore, Good is in a better epistemic position than Bad to satisfy the proper basing condition for doxastic justification.

Good reasoning is crucial for satisfying both ideal and non-ideal standards of rationality. If Good and Bad are non-ideal agents, then neither is capable of reasoning well enough to satisfy ideal standards of rationality. Recall Christensen's example: each of them has propositional justification to be certain that the theorem is true and hence to bet the house against a nickel in the certainty of winning. But neither can be certain of winning in a way that is doxastically justified, since neither has doxastic dispositions that are perfectly sensitive to the distinction between theorems and non-theorems. Even so, there may be some high but less than maximal degree of confidence that is required by ordinary standards of rationality. And Good is in a better position than Bad to satisfy those ordinary standards of rationality since his doxastic dispositions are more sensitive to the distinction between theorems and non-theorems.

On this view, the epistemic role of experience in the apriori domain is an *enabling* role, rather than a *justifying* role: it is not a source of propositional justification to believe logical truths, but it plays an important role in enabling subjects to use the propositional justification that they already have in forming doxastically justified beliefs. The distinction between enabling and justifying roles for experience is a familiar one: for instance, it is a routine point that color experience plays a role in enabling subjects to understand the proposition that

nothing can be red and green all over, without thereby providing the source of their justification to believe it.[21] But my proposal goes beyond this in two ways. First, the enabling role of experience encompasses not just perceptual experience, but also cognitive experience, including experiences of understanding, reasoning, and intuition. And second, the role of experience is not restricted to its role in enabling the grasp of propositions, but includes whatever else is needed to convert propositional justification into doxastic justification. The enabling role of experience includes much more on this view than is traditionally recognized.[22]

7. The Disabling Role of Experience

Christensen (2007) argues that the requirement of logical omniscience is inconsistent with the requirement to respect empirical evidence of one's own cognitive imperfection. He gives the following example to illustrate the tension:

> Suppose I work out my proof of T after having coffee with my friend Jocko. Palms sweaty with the excitement of logical progress, I check my work several times, and decide that the proof is good. But then a trusted colleague walks in and tells me that Jocko has been surreptitiously slipping a reason-distorting drug into people's coffee – a drug whose effects include a strong propensity to reasoning errors in 99% of those who have been dosed (1% of the population happen to be immune). He tells me that those who have been impaired do not notice any difficulties with their own cognition – they just make mistakes; indeed, the only change most of them notice is unusually sweaty palms. (2007: 10)

Christensen argues that it is rational to be certain that T is true only if it is also rational to be certain that E is false, where E is an error proposition that is

---

[21] See Burge (1993: 460), BonJour (1998: 9-10) and Chalmers (2012: 189-90).
[22] See Peacocke (2004: 205-7), Chalmers (2012: 185-98) and Ichikawa and Jarvis (2013: 166-70) for similarly expansive conceptions of the enabling role of experience, although none of them goes so far as to endorse the claims about the disabling role of experience proposed in sections 7 and 8.

inconsistent with T: *I falsely believe T on the basis of mistaken reasoning*. Moreover, he continues, it is not rational to be certain that E is false given strong evidence that E is true in the form of testimony from a trusted colleague corroborated by an experience of unusually sweaty palms. Therefore, he concludes, it is not rational to be certain that T is true and so rationality does not require logical omniscience.

Notoriously, of course, one philosopher's modus tollens is another philosopher's modus ponens. So one might insist that since it is rational to be certain that T is true, it is also rational to be certain that E is false. But Christensen has a compelling reply: it would be irrational and dogmatic for you or me to disregard evidence of our cognitive imperfections on the basis of the following Certainty Argument:

(1) It is certain that T is true.
(2) If it is certain that T is true, then it is certain E is false.
(3) Therefore, it is certain that E is false.

I agree with Christensen that there is something wrong with using the Certainty Argument in this way, but I disagree about what exactly is wrong with it.

Christensen's view is that evidence of one's cognitive imperfection functions as a *propositional defeater* – that is to say, it defeats one's propositional justification to be certain of logical truths.[23] In the Jocko case, for example, one's evidence that E is true defeats one's propositional justification to be certain that T is true. On this view, what's wrong with the Certainty Argument is that evidence that justifies doubts about its conclusion thereby defeats one's justification to be certain of its premises. I cannot accept this account since I maintain that apriori logical justification is indefeasible: one's apriori justification to be certain that T is true defeats one's aposteriori justification for positive credence that E is true, rather than

---

[23] See also Feldman (2005), Christensen (2010), and Kelly (2010), although all of these authors acknowledge that evidence of one's cognitive imperfection is in some respects different from more standard examples of undercutting defeaters.

vice versa. So the challenge for me is to give an alternative account of what's wrong with using the Certainty Argument.

I claim that evidence of one's cognitive imperfection functions as a *doxastic defeater*, rather than a *propositional defeater* – that is to say, it does not defeat one's propositional justification to be certain that T is true, and hence that E is false, but merely prevents one from converting this propositional justification into doxastic justification. On this view, what's wrong with using the Certainty Argument is that although one has propositional justification to be certain that T is true, and hence that E is false, one is nevertheless disabled from being certain in a way that is doxastically justified. As a result, one cannot use the Certainty Argument to assuage doubts about its conclusion.[24]

On this account, the epistemic role of experience in the logical domain is not a justifying and defeating role, but rather an enabling and disabling role. Just as experience plays a role in enabling us to convert propositional justification into doxastic justification, so it plays a role in disabling us from doing so. But this proposal about the role of experience applies only to non-ideally rational agents like you and me. (I defer discussion of ideally rational agents until section 8.)

Empirical evidence of one's cognitive imperfection has a disabling effect on non-ideally rational agents because their doxastic dispositions are not sufficiently sensitive to the logical facts. If I use the Certainty Argument in becoming certain that E is false, then my belief is doxastically unjustified since it is held on the basis of doxastic dispositions that are not appropriately sensitive to the distinction between "the good case" in which I correctly prove a theorem and "the bad case" in which I make a mistake. Exercising the same doxastic dispositions in response to the same empirical evidence could easily yield the false and unjustified belief that I correctly proved the theorem when in fact I didn't. As a result, I do not satisfy the proper basing condition for doxastic justification as explained in section 6.

---

[24] Compare Pryor's (2004) account of what's wrong with Moore's argument for the existence of the external world. According to Pryor, there is nothing wrong with the justificatory structure of the argument, although there are limits on its dialectical effectiveness: "Anyone who had doubts about its conclusion couldn't *use* the argument to rationally overcome those doubts" (2004: 363).

We can distinguish two aspects of the function of empirical evidence of one's own cognitive imperfection. On the one hand, it disables non-ideal agents from complying with ideal standards of rationality. On the other hand, it has an effect on what is required for complying with non-ideal standards of rationality. Even if I cannot be certain of logical truths in a way that meets ideal standards of rationality, I can be more or less confident of such truths in ways that meet ordinary standards of rationality. But evidence of one's cognitive imperfection impacts on how much confidence is required by ordinary standards of rationality.

In particular, evidence that I've ingested a reason-distorting drug increases the degree of confidence that is non-ideally reasonable to invest in E and thereby reduces the degree of confidence that is non-ideally reasonable to invest in T. Two non-ideal agents might be equally reliable in reasoning about logic, but if one of them acquires evidence that she has ingested a reason-distorting drug, then it is reasonable by non-ideal standards for her to invest less confidence in T. How much confidence it is non-ideally reasonable to invest in logical truths depends not just on one's degree of cognitive imperfection but also on the state of one's evidence about one's cognitive imperfection.[25]

How does empirical evidence of one's cognitive imperfection affect which attitudes are required by non-ideal standards of rationality? In section 5, I proposed that the option required by non-ideal standards of rationality is the best option in the ideal ranking that the agent can take on the basis of sufficiently good reasoning. But misleading evidence about your cognitive imperfection impacts neither the actual quality of our first-order reasoning nor the ideal ranking of options. So there is something of a puzzle here.

My response is that acquiring evidence about our cognitive imperfection brings new reasoning dispositions into play and so it does impact on the overall quality of our reasoning. We need to consider not just your first-order dispositions

---

[25] This account of the function of evidence of one's cognitive imperfection has more general implications for debates about the epistemic significance of disagreement, but I do not have space to explore them here. See van Wietmarschen (2013) for an account that is broadly congenial to mine.

to reason about logic, but also your second-order dispositions to respond to evidence of your own cognitive imperfection. Your first-order reasoning is reliable enough to make high credence rational in ordinary contexts, but your second-order dispositions to respond to evidence of your cognitive imperfection are unreliable enough to make high credence irrational in contexts where you have such evidence. That is why misleading evidence of one's cognitive imperfection can reduce the credence that is required by ordinary standards of rationality.

8. Ideally Rational Agents

The account of the disabling role of experience proposed in section 7 does not extend from non-ideally rational agents (NRAs) to ideally rational agents (IRAs). Empirical evidence of one's cognitive imperfection plays a disabling role for NRAs because our logical reasoning is not good enough: our doxastic dispositions are not sufficiently sensitive to the logical facts. This is why it seems irrational and dogmatic for non-ideal agents like us to use the Certainty Argument because we're disposed to use the same reasoning for logical truths and falsehoods alike.

IRAs are immune from these disabling effects, since their logical reasoning is perfectly sensitive to the logical facts. This is not to deny that an IRA can acquire misleading empirical evidence that she is cognitively imperfect. Nor is it to claim that she can simply ignore it. The claim is rather that an IRA can use the Certainty Argument in coming to know with certainty that an error proposition is false when it is inconsistent with the apriori truths of logic. An IRA, unlike an NRA, is not disposed to use the same reasoning for logical truths and falsehoods alike because her logical reasoning is perfectly sensitive to the logical facts.

Does this mean IRAs are immune to the effects of reason-distorting drugs? This would be surprising, since having an iron constitution is not a plausible condition for ideal rationality. Better to say that IRAs can be drugged, but the effect of these drugs is to change their doxastic dispositions and thereby to make them non-ideal. We needn't accept the principle, "Once an IRA, always an IRA." One could be ideal at one time and non-ideal at another. Indeed, one could be ideal at the same time that one is in danger of becoming non-ideal. The danger of becoming non-ideal

does not thereby make one non-ideal any more than the danger of becoming insane thereby makes one insane. One's first-order reasoning dispositions may be perfectly sensitive to the logical facts even if one has a second-order disposition to acquire first-order reasoning dispositions that are imperfectly sensitive – say, because one is in danger of ingesting Christensen's reason-distorting drugs or inhaling Hawthorne's "apriori gas".[26]

Does this mean IRAs are certain of their own cognitive perfection? This would be surprising too, since empirical knowledge about one's own reasoning dispositions is not a plausible condition for ideal rationality. Christensen puts the objection as follows:

> An IRA, as usually conceived of in theorizing about rationality, is quite different from an omniscient god. The IRA reasons perfectly, and is thus logically omniscient…, but the IRA is not assumed to be factually omniscient. This conception of an IRA carries with it no obvious presumption that an IRA would *know* that she was ideally rational. (2007: 14)

IRAs need not be certain of their own cognitive perfection at past or future times, since an agent can be ideal at one time and non-ideal at another. Moreover, an IRA need not be rationally certain of her own cognitive perfection at the present time, so long as she is rationally certain of any logical truth T and hence certain of the falsehood of any error proposition E that is inconsistent with T.

Christensen observes that there is a bootstrapping problem here. If an IRA can be rationally certain that T is true, and hence that E is false, then she can repeat this performance for a whole series of logical truths. Indeed, if an IRA is cognitively unlimited, then there is no limit to the number of logical truths for which she can repeat this performance; in principle, she can do it for all of them. And so, if an IRA

---

[26] Hawthorne (2007) uses the example of "apriori gas" to argue that apriori knowledge does not supervene on one's intrinsic properties given a safety condition for knowledge. However, the safety condition is not violated in the presence of apriori gas so long as it is formulated in a basis-sensitive way.

can be rationally certain about what her doxastic attitudes are at any given time, then she can also become rationally certain that she satisfies the requirements of ideal rationality: for any logical truth, she is certain that it is true. Doesn't it follow that an IRA must be certain of her own cognitive perfection?

Not quite. An IRA does not just hold all the doxastic attitudes that she has propositional justification to hold; she also holds them in a way that is doxastically justified on the basis of perfectly good reasoning. Even if an IRA can be certain that she satisfies the first condition, she cannot be certain that she satisfies the second. At best, the Certainty Argument delivers the conclusion that, for every logical truth, she is certain that it is true and so she holds the doxastic attitude that she has propositional justification to hold. It does not deliver the conclusion that her doxastic attitude is formed on the basis of perfectly good reasoning. An IRA must be certain of the facts about propositional justification, but she need not be certain of the facts about doxastic justification, since she can have rational doubts about the reasoning on the basis of which her beliefs are held.[27]

The upshot is that ideal rationality is consistent with a certain kind of Moorean predicament. Consider the pair of Moorean conjunctions below:

> *Malignant*: T and I don't have [propositional] justification to believe that T.
> *Benign*: T and I don't believe that T in a way that is [doxastically] justified.

I have argued (in Smithies 2012c) that one cannot have evidence that justifies believing Moorean conjunctions of the *malignant* form. But one can certainly have evidence that justifies believing Moorean conjunctions of the *benign* form. Suppose one has empirical evidence that a demon has debased one's beliefs without disturbing the match between one's beliefs and the justifying evidence.[28] In that case, there is no rational pressure to revise one's beliefs.

---

[27] This fits with a version of access internalism on which propositional justification but not doxastic justification is accessible by reflection alone. See Smithies (2012c).
[28] See Schaffer (2010) for the debasing demon.

An IRA can be in just this kind of benign Moorean predicament. Suppose she has convincing empirical evidence that she has taken a reason-distorting drug that causes reasoning errors in the majority of the population and that she is not immune from its effects. The most probable hypothesis given her total evidence is perhaps that she made an error in reasoning that leads to a true and certain conclusion. In that case, rationality requires that she doubt the quality of her reasoning, but since she is certain that her conclusion is true, there is no rational pressure to change her view. So, she is required to believe the benign Moorean conjunction, "It is certain that T, but I am not certain that T in a way that is justified."

9. Rational Dilemmas

The arguments of this paper have more general implications for the existence of rational dilemmas. Christensen argues that evidence about one's own cognitive imperfection give rise to a tension between the following rational ideals:

(1) (LOGIC) An agent's beliefs must respect logic by satisfying (some version of) probabilistic coherence.
(2) (EVIDENCE) An agent's beliefs (at least about logically contingent matters) must be proportioned to the agent's evidence.
(3) (INTEGRATION) An agent's object-level beliefs must reflect the agent's meta-level beliefs about the reliability of the cognitive processes underlying her object-level beliefs. (2007: 20)

The idea is that an agent whose first-order beliefs respect (LOGIC) but whose second-order beliefs about her own cognitive imperfections respect (EVIDENCE) thereby fails to exhibit the kind of inter-level coherence that is required by (INTEGRATION). So, it is impossible to satisfy all three ideals given evidence of one's own cognitive imperfection. Christensen articulates the point as an epistemic analogue of Murphy's Law:

The usual version of Murphy's Law states that if it's possible for something to go wrong, it will. The epistemic cousin says that if it's possible that something has gone epistemically wrong (more specifically, if it's possible that I've made a mistake in thinking about some theorem T), then something has actually gone epistemically wrong (my belief about T falls short of some rational ideal). For either I'm certain of T, in which case my belief fails to reflect appropriately the possibility that I've made a cognitive error, or I'm uncertain about T, in which case my belief fails to respect logic. (2007: 24-5)

Nevertheless, Christensen argues, it would be a mistake to think that these are not all genuine rational ideals. Here is Christensen's diagnosis:

We're quite familiar with other ideals that operate as values to be maximized, yet whose maximization must in certain cases be balanced against, or otherwise constrained by, other values…. [T]he fact that ideals can be in tension with one another does not undermine their status as ideals. So we can still see (LOGIC) as a rational ideal once we see how it is to be constrained by (EVIDENCE) and (INTEGRATION). (2007: 24)

On this view, a probabilistically coherent agent would not be ideally rational, since this means violating other rational ideals – namely, (EVIDENCE) or (INTEGRATION). Nevertheless, it can be useful to think of a probabilistically coherent agent as maximizing one dimension of ideal rationality that needs to be weighed against others.[29] On this view, there are a number of distinct rational ideals, each of which captures some dimension of ideal rationality, but none of which can be maximized without thereby sacrificing others. So, on this view, there cannot be an agent that maximizes all dimensions of ideal rationality; the best we can hope for from an

---

[29] Compare Chalmers's (2012: 101-7) response to the problem of self-doubt by postulating *ideal insulated cognizers*: "cognizers whose rational processes are practically insulated from higher-order beliefs…but otherwise ideal."

ideally rational agent is some maximally good way of weighing these competing dimensions against one another.[30]

My arguments in this paper make available an alternative diagnosis on which there is no inherent tension between these dimensions of ideal rationality. On this diagnosis, Murphy's Law applies to non-ideal agents, but not ideally rational agents. Ideally rational agents are capable of maximizing (LOGIC), (EVIDENCE), and (INTEGRATION), whereas non-ideally rational agents are incapable of doing so. On this view, the tension between these various dimensions of ideal rationality is not inherent in the structure of ideal rationality itself, but is rather a reflection of the plight of non-ideal agents. Given our fallen state, we are forced to trade off certain dimensions of ideal rationality against others and there may be some degree of incommensurability in determining which trade offs are best for us to make. But this is no reflection of the structure of ideal rationality itself.

10. Conclusions

In this paper, I have argued for the following claims. First, a negative thesis about the source of apriori logical justification: it does not have its source in facts about one's experience. Second, a positive thesis about the source of apriori logical justification: it has its source in the logical facts that one has apriori justification to believe. Third, a thesis about the enabling role of experience in the logical domain: it plays a role in enabling subjects to convert propositional justification into doxastic justification. Fourth, a thesis about the disabling role of experience in the logical domain: it plays a role in disabling subjects from converting propositional justification into doxastic justification. Fifth, a thesis about ideally rational agents: they are immune from the disabling role of experience in the logical domain. Sixth, a thesis about the structure of rationality: there are no rational dilemmas that result from conflicts between first-order evidence and higher-order evidence. Finally, the

---

[30] See Lasonen-Aarnio (2014) for criticism of this proposal.

33

paper as a whole provides a strategy for motivating and defending the thesis that ideal rationality requires omniscience and infallibility about logic.[31]

References

Alston, W. 1971. Varieties of Privileged Access. *American Philosophical Quarterly* 8.3: 223-41.

Alston, W. 1988. The Deontological Conception of Epistemic Justification. *Philosophical Perspectives* 2: 257-99.

Bengson, J. Forthcoming. The Intellectual Given. *Mind*.

Boghossian, P. 2000. Knowledge of Logic. In *New Essays on the Apriori*, edited by P. Boghossian & C. Peacocke. Oxford University Press.

BonJour, L. 1998. *In Defense of Pure Reason*. Cambridge University Press.

Burge, T. 1993. Content Preservation. *Philosophical Review* 102.4: 457-88.

Burge, T. 1998. Frege on Knowing the Foundation. *Mind* 107: 305-47.

Casullo, A. 2003. *A Priori Justification*. Oxford University Press.

Chalmers, D. 2011. The Nature of Epistemic Space. In *Epistemic Modality*, edited by A. Egan & B. Weatherson. Oxford University Press.

Chalmers, D. 2012. *Constructing the World*. Oxford University Press.

Christensen, D. 2004. *Putting Logic in Its Place*. Oxford University Press.

Christensen, D. 2007. Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals. *Oxford Studies in Epistemology* 2: 3-31.

Christensen, D. 2010. Higher-Order Evidence. *Philosophy and Phenomenological Research* 81.1: 185-215.

Chudnoff, E. 2013. *Intuition*. Oxford University Press.

Feldman, R. and Conee, E. 1985. Evidentialism. *Philosophical Studies* 48.1: 15-34.

Feldman, R. 2000. The Ethics of Belief. *Philosophy and Phenomenological Research* 60.3: 667-695.

Feldman, R. 2005. Respecting the Evidence. *Philosophical Perspectives* 19: 95-119.

Foley, R. 1993. *Working Without a Net: A Study of Egocentric Epistemology*. Oxford University Press.

Hajek, A. MS. Staying Regular.

Hawthorne, J. 2002. Deeply Contingent A Priori Knowledge. *Philosophy and Phenomenological Research* 65.2: 247-69.

Hawthorne, J. 2007. Apriority and Externalism. In *Internalism and Externalism in Semantics and Epistemology*, edited by S. Goldberg. Oxford University Press.

Huemer, M. 2001. *Skepticism and the Veil of Perception*. Rowman and Littlefield.

Ichikawa, J. and Jarvis, B. 2013. *The Rules of Thought*. Oxford University Press.

Kahneman, D., Slovic, T., and Tversky, A. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

Kelly, T. 2010. Peer Disagreement and Higher-Order Evidence. In *Disagreement*, edited by R. Feldman & T. Warfield. Oxford University Press.

Kratzer, A. 1981. The Notional Category of Modality. In *Words, Worlds, and Contexts*, edited by H. Eikmeyer and H. Rieser. De Gruyter.

Kvanvig, J. and Menzel, C. 1990. The Basic Notion of Justification. *Philosophical Studies* 59: 235-61.

Lasonen Aarnio, M. 2014. Higher-Order Evidence and the Limits of Defeat. *Philosophy and Phenomenological Research* 88.2: 314-45.

Peacocke, C. 2004. *The Realm of Reason*. Oxford University Press.

Pryor, J. 2001. Highlights of Recent Epistemology. *British Journal for the Philosophy of Science* 52: 95-124.

Pryor, J. 2004. What's Wrong With Moore's Argument? *Philosophical Issues* 14.1: 349-78.

Schaffer, J. 2010. The Debasing Demon. *Analysis* 70.2: 228-37.

Schoenfield, M. 2012. Chilling Out On Epistemic Rationality. *Philosophical Studies* 158.2: 197-219.

Schoenfield, M. MS. Bridging Rationality and Accuracy.

Smithies, D. 2012a. The Simple Theory of Introspection. In D. Smithies and D. Stoljar, *Introspection and Consciousness*. Oxford University Press.

Smithies, D. 2012b. Mentalism and Epistemic Transparency. *Australasian Journal of Philosophy* 90.4: 723-41.

Smithies, D. 2012c. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85.2: 273-300.

Sosa, E. 2003. Privileged Access. In *Consciousness: New Philosophical Perspectives*, edited by Q. Smith & A. Jokic. Oxford University Press.

Titelbaum, M. Forthcoming. Rationality's Fixed Point. *Oxford Studies in Epistemology*.

Turri, J. 2010. On the Relationship Between Propositional and Doxastic Justification. *Philosophy and Phenomenological Research* 80.2: 312-26.

Van Wietmarschen, H. 2013. Peer Disagreement, Evidence, and Well Groundedness. *Philosophical Review* 122.3: 395-425.

Watson, G. 1975. Free Agency. *The Journal of Philosophy* 72: 205-20.

Williamson, T. 2007. On Being Justified in One's Head. In M. Timmons, J. Greco, and A. Mele, *Rationality and the Good*. Oxford University Press.