

Inferential Justification and the Transparency of Belief

Abstract. This paper critically examines currently influential transparency accounts of our knowledge of our own beliefs that say that self-ascriptions of belief typically are arrived at by “looking outward” onto the world. For example, one version of the transparency account says that one self-ascribes beliefs via an inference from a premise to the conclusion that one believes that premise. This rule of inference reliably yields accurate self-ascriptions because you cannot infer a conclusion from a premise without believing the premise, and so you cannot infer from a premise that you believe the premise unless you do believe it. I argue that this procedure cannot be a source of justification, however, because one can be justified in inferring from p that q only if p amounts to strong evidence that q is true. This is incompatible with the transparency account because p often is not very strong evidence that you believe that p . For example, unless you are a weather expert, the fact that it *will* rain is not very strong evidence that you *believe* it will rain. After showing how this intuitive problem can be made precise, I conclude with a broader lesson about the nature of inferential justification: that inferential transitions between beliefs, when justified, must be underwritten by evidential relationships between the facts or propositions which those beliefs represent.

Keywords. transparency, inference, evidence, Moore’s paradox, introspection

Word Count. 12,177 (excluding notes and references)

1. Introduction

You can know what you believe without appealing to the sort of evidence you would need in order to know what another person believes. For example, you might know that you believe that there will be a third world war, even though you have never heard yourself say so out loud. And you might know that you believe the capital of North Dakota is Bismark, even though nothing about your recent behavior would distinguish you from someone who has forgotten what the capital of North Dakota is. How do you know these things about yourself?

One familiar proposal is that you know your own beliefs by deploying an ‘inner sense’—a faculty of perception somehow comparable to ordinary perceptual faculties, but directed inward rather than outward. According to a number of recent theorists, however, it is a mistake to invoke the notion of perception to explain your knowledge of your own beliefs. To answer the question ‘Do I believe that p ?’, these theorists claim, there is no need to “look inward” to examine or observe your own mind. Instead you need only “look outward” onto the world, and answer the question ‘Is it the case that p ?’. If you

look outward to find that it will rain, for example, then you can self-ascribe the belief that it will rain on the basis of what you have found.¹ In the form of a common slogan, the proposal is that questions about one's own beliefs are *transparent* to questions about the mind-independent world.²

Although many theorists are attracted to this rough and suggestive slogan, not everyone who likes the slogan understands its meaning in quite the same way. For simplicity, we can focus on a specific commitment shared by a number of views whose differences are for our purposes inessential. Let's say (stipulatively) that *p* is a **good reason** for you to believe that *q* if you either already are in a position justifiably to believe that *q* on the basis that *p*, or if you would be in such a position if you were to learn that *p*. For example, if you know that the barometer in the garage is working properly, then the fact that the barometer reads 'low' will be a good reason for you to believe that the barometric pressure is low in the absence of further evidence to the contrary. We can now represent the claim that questions about one's beliefs are transparent to questions about the world schematically as follows:

Transparency: *p* is a good reason for you to believe that you believe that *p*.³

¹ Related proposals have been suggested concerning knowledge of mental states other than beliefs. For proposals concerning knowledge of one's own perceptual experiences, see (Byrne, 2010), (Dretske, 1994 and 1995, Ch. 2), (Evans, 1982), and (Tye, 2002), and for critical discussion see (Aydede, 2002) and (Lycan, 1999). For proposals concerning desire and intention, see (Byrne, 2011 and forthcoming a), and (Setiya, 2011). For epistemic states like justified belief and knowledge, see (Gibbons, 2006, pp. 32-33) and (McHugh, 2010). It is controversial whether we should expect an account of self-knowledge to be uniform across these cases, and it may be that a transparency-based explanation is correct in some cases but not others. For discussion, see (Boyle, 2009 and 2011), (Byrne, 2011, Sec. 5 and forthcoming a, Sec. 3), (Cassam, 2011), and (Schwitzgebel, forthcoming).

² See, e.g., (Cassam, 2011), (Moran, 2001), (Shah and Velleman, 2005). Other theorists express a similar idea by saying that beliefs are transparent to the world. Although the term of art 'transparency' often is associated with Evans' influential discussion of belief self-ascription (1983, pp. 225-227), Evans himself did not use the term. Sometimes 'transparency' is used instead to denote an (alleged) property of experiences (a.k.a., 'diaphanousness'). See, e.g., (Stoljar, 2004) and (Tye, 2002).

³ With minor, largely terminological differences, Transparency is endorsed by Byrne (2005, Sec. 7), Gallois (1996), and Shoemaker (2009, pp. 37-39)—although Byrne's focus is knowledge rather than justification. Other attributions are less straightforward. Moran (2001 pp. 62-63), offers two formulations of what he calls "the claim of transparency" that are not obviously equivalent to one another. The first, following Edgley (1969, p. 90), is that "from within the first-person perspective, I treat the question of my belief about *P* as equivalent to the question of the truth of *P*". The second, which follows Evans, is that "a first-person present-tense question about one's belief is answered by reference to (or consideration of) the same reasons that would justify an answer to the corresponding question about the world." I think both of these proposals face difficulties along the lines of those I discuss here, but to avoid redundancy I will not argue the point. For further examples of the proposal that you can be justified in self-ascribing the belief that *p* on the basis of *evidence* for *p*, rather than on the basis of *p* itself, see (Fernández, 2003 and 2005) and (Williams, 2007), and for critical discussion see (Vahid, 2005), (Way, 2007), and (Zimmerman, 2004). Setiya (2011) and Silins (2012) disavow what I call 'Transparency' in favor of closely related alternatives, while Boyle (2011) favors an alternative view whose relationship to Transparency is more distant. For discussion of (Peacocke, 1998), whose account of consciously based self-ascriptions is a precursor of recent transparency accounts, see (Coliva, 2008), (Martin, 1998), and (Zimmerman, 2006).

Despite its attractions, which we will consider shortly, Transparency can appear puzzling because it does not fit with our ordinary understanding of what a good reason (or basis) for belief is. To a first approximation, we usually think that p is a good reason for you to believe that q only if p amounts to strong *evidence* that q is true. This conception of epistemic reasons does not sit well with Transparency, because p often is not very strong evidence that you believe that p . For example, unless you are some kind of hyper-expert about the weather, who can be counted on not to be ignorant of whether it will rain, the fact that it *will* rain is not very strong evidence that you *believe* it will rain.

Transparency has been gaining traction in recent discussions of self-knowledge, and this evidentialist objection, which this paper will develop in greater detail, is arguably its most pressing difficulty.⁴ But there is a further reason why this objection to Transparency deserves of our attention. As we will see, it lies at the fault line between competing conceptions of inferential justification, one of which gives explanatory priority to inferential relations between one's beliefs, and the other of which gives priority to evidential relations between the objects of those beliefs. So it is no small matter whether Transparency, and the corresponding conception of inferential justification, can be sustained in the face of its most pressing difficulty.

After outlining some of the key motivations for Transparency in Section 2, in Section 3 we will set aside a worry to which the Transparency theorist may have an effective reply, and take away a broader lesson about a theory of inferential justification which supports Transparency. Then in Section 4, we will consider a series of specific difficulties for Transparency, which I take to clarify and reinforce the intuitive puzzle. In Section 5 we will take a step back and consider a broader lesson about the nature of inferential justification: that inferential transitions between beliefs, when justified, must be underwritten by evidential relationships between the facts or propositions which those beliefs represent.

2. Motivations for Transparency

One motivation for Transparency is that it fits neatly with a popular idea about the relationship between rationality and self-knowledge. As G. E. Moore famously observed, it seems somehow inappropriate (or “absurd”) to *assert* propositions of the form $\langle p, \text{ but I don't believe that } p \rangle$.⁵ Many subsequent theorists have thought we should go further, and countenance as irrational combinations of attitudes with a broadly Moore-paradoxical character. More specifically, these latter-day Moorean theorists have urged us to accept claims like the following:

⁴ See, e.g., (Boyle, 2011), (Gallois, 1996), (Martin, 1998, pp. 117-118), (Moran, 2003, p. 413), (Silins, 2012), and especially Byrne's (2005, Sec. 7.2) discussion of “the puzzle of transparency.” We will consider some of these discussions in Section 4.

⁵ For a helpful review, see (Green and Williams, 2007).

Moorean Irrationality: (1) It is irrational to believe that you believe that p if you do not believe that p, and (2) It is irrational to doubt⁶ that you believe that p if you do believe that p.

To better understand Moorean Irrationality, it will help to consider what its proponents allege is an important difference between introspective knowledge of your own beliefs and proprioceptive knowledge of your bodily states such as the state of having your legs crossed. According to the Moorean, the state of jointly believing that p and yet doubting that you believe that p manifests a kind of internal incoherence among your doxastic attitudes which is allegedly the mark of irrationality.⁷ In contrast, the state of having your legs crossed and yet doubting that you have your legs crossed does not manifest an incoherence among your attitudes for the simple reason that the state of having your legs crossed is not a doxastic attitude.

To see why Moorean Irrationality is a close cousin of Transparency, we need to introduce a bit of jargon. Let's say, roughly and approximately, that believing that p **rationally requires** you to believe that q if you are in a situation where it is irrational for you to doubt that q if you believe that p, but where it could be rational for you to doubt that q if you doubt that p.⁸ It arguably follows from Moorean Irrationality that

Moorean Requirements: Believing that p rationally requires you to believe that you believe that p.⁹

There is a close affinity between Moorean Requirements and Transparency.¹⁰ When we consider paradigmatic examples in which the belief that p rationally requires the belief that q, it also is true that p is a good reason to believe that q—that one could be justified in believing that q on the basis that p. For example, given the right background evidence the belief that it will rain might rationally require you to believe that a tennis match will be cancelled. And given the very same background evidence, it seems that you would then be justified in believing that the match will be cancelled on the basis that it will rain.

⁶ I take **doubting** that p to mean either disbelieving or actively withholding belief from p. Someone who fails to believe that p merely because they have never considered whether p does not doubt that p in this sense.

⁷ See, e.g., (Christensen, 2007) and (Zimmerman, 2008).

⁸ This first pass gloss on rational requirements is complicated by cases where one's attitude to p is itself irrational. We will consider such cases in greater detail in Section 4.2 below.

⁹ Moorean Requirements follows from Moorean Irrationality on the plausible assumption that if it is irrational for you to believe that q then it is rational for you to doubt that q. See the discussion of the Existence thesis in Barnett (MS a) for more.

¹⁰ For related discussion, see (Fernández, 2005a), (Moran, 2001, pp. 69-77), (Shoemaker, 2009, pp. 37-39), (Silins, 2012), (Smithies, 2012 and forthcoming), (Williams, 2007), and (Zimmerman, 2008).

Nevertheless, the Moorean might have room for a fallback position that avoids Transparency by distinguishing between propositions that serve as good reasons for a given belief and propositions which if believed would rationally require that belief. One strategy for doing so is to distinguish rational requirement from justification, and to say that what rationally requires a belief need not in doing so justify that belief.¹¹ This move could avoid the implication from Moorean Requirements to Transparency because it allows us to say that belief that *p* rationally requires belief that one believes that *p* without justifying belief that one believes that *p*. But it would do so at the cost of severing the explanation of how self-ascriptions are justified from considerations involving Moore's paradox in a way that some Mooreans might find unsatisfying. Another strategy is to distinguish instead between two ways in which believing that *p* might justify believing that *q*, a reason-providing way and another as yet unexplained way. When believing that *p* justifies believing that *q* in the typical reason-providing way, the fact or proposition that *p* is a good reason or basis on which to believe that *q*, and one's believing that *p* justifies believing that *q* by providing one with that reason. But perhaps there is some other way that the belief that *p* might confer justification onto the belief that one believes that *p*—a way that does not involve providing one with reasons on whose basis one can believe that one believes that *p*.¹² This strategy avoids whatever difficulties Transparency might face, but only by offering us little positive insight into the nature of introspective justification. Once it is conceded that the belief that *p* does not justify the belief that one believes *p* in the way that beliefs usually can confer justification on other beliefs, it is unclear what explanatory value remains to the claim that beliefs justify their own self-ascriptions, for it is unclear what grip we have on the notion of a belief justifying another belief apart from our grip on the conventional reason-providing kind of justifying. To be sure, many philosophers have proposed views under which mental states other than beliefs—such as perceptual experiences—can justify beliefs, just as other beliefs can. If we are open to the possibility of non-beliefs justifying beliefs, then perhaps we also should be open to beliefs justifying other beliefs in unconventional ways. However, the most prevalent and plausible theories of perceptual justification are ones that try to assimilate justification by perceptual experience to the conventional, reason-giving sort of justification that usually comes from beliefs, and show us how these can be two related species of a common genus.¹³ It is possible that the same could be done by the Moorean, but in the absence of

¹¹ Avnur (2012) adopts a position like this, but in a very different context. On Avnur's view, it is possible for our ordinary beliefs to rationally require beliefs in the negations of skeptical hypotheses without thereby justifying those beliefs.

¹² I take this, in outline, to be Silins' (2012) strategy, although on his view it is the judgement that *p*, rather than the belief that *p*, that justifies the belief that one believes that *p*. Although Silins characterizes his view as an elaboration of the slogan that the question whether you believe that *p* is transparent to the question whether *p*, on his construal the slogan is consistent with a wide variety of theories about belief-self ascription, including an inner sense theory. Jim Pryor and Declan Smithies have suggested similar ideas to me in conversation.

¹³ See Pryor (2005) for a critical discussion of attempts to assimilate perceptual justification to justification by beliefs via what Pryor calls the 'Premise Principle'.

a convincing account we should not take this for granted. Just because the phenomenon of reason-providing justification between one belief and another is one we have a grip on, we cannot take it for granted that we also have a grip on what it would mean for one belief to justify another in the unconventional way that the Moorean proposes.

A second motivation for Transparency draws on a theory about the psychological process underlying typical self-ascriptions of belief. Where an **inference** is understood to be the occurrent process of reasoning by which one forms the standing attitude of believing a conclusion on the basis of a reason or premise, Alex Byrne recently has proposed the striking theory that

Transparent Inferences: Typical psychological transitions from a belief that *p* to a belief that one believes that *p* are inferences.¹⁴

Byrne's guiding idea is the observation that ordinary reasoning and the introspective self-ascription of belief have an important feature in common: namely, both involve forming new beliefs in a manner that is sensitive to one's preexisting beliefs. Since a capacity for reasoning and inference is necessary to explain much of our knowledge of the external world, Transparent Inferences is alleged to provide an *economical* explanation of the reliability of our ordinary self-ascriptions of belief—an explanation which, in contrast to the allegedly *extravagant* inner sense theory, posits no psychological mechanisms for detecting one's beliefs beyond those required for psychological capacities other than introspection.

It is natural to think that if we accept Transparent Inferences then we also must accept Transparency. Since our typical self-ascriptions of belief are justified, we will be committed to Transparency if we accept the descriptive claim that we do typically infer from *p* that we believe that *p* and yet deny Transparency. But there might be room for a fallback position that rejects this descriptive claim in favor of the closely related claim that

Redeployment of Reasoning: Typical psychological transitions from a belief that *p* to a belief that one believes that *p* redeploy the psychological mechanisms ordinarily deployed in making inferences.

Unlike Transparent Inferences, which concerns the reasons or basis on which you typically self-ascribe beliefs, Redeployment of Reasoning speaks only of the the psychological mechanisms that are responsible for generating self-ascriptions. These kinds of claims arguably can come apart in cases where psychological mechanisms

¹⁴ See especially (Byrne, 2005, Sec. 7). Byrne's proposal was anticipated by Gaillois' (1996) doxastic schema, and by a distinct but closely related proposal from Fernández (2003, 2005b). Despite some important differences, Peacocke's (1998) account of consciously based self-ascriptions might be interpreted along the same lines.

typically deployed to serve one function are deployed to serve a different function, as the Redeployment of Reasoning theorist claims for cases of belief self-ascription.

Although this fallback position is able to avoid commitment to Transparency and the epistemic consequences that come with it, it does so only by severing the epistemology of self-ascriptions from their psychological explanation. This will be unsatisfying to many theorists, who wish to connect Moorean intuitions about the epistemology of self-ascriptions to an account of the psychological process by which we arrive at them.

What is more, the alleged psychological economy of the Redeployment of Reasoning hypothesis is a thin empirical basis for its acceptance. Even if we accept explanatory economy as a desiderata for psychological explanations, Redeployment of Reasoning is not the only game in town when it comes to redeploying for the sake of belief self-ascription cognitive mechanisms that ordinarily subserve other functions. For just one alternative, consider the Ryle-inspired theory that we know what we think in part by “hearing” our own “silent soliloquies.”¹⁵ On a theory like this, we often know what we believe by the production and monitoring of speech-like mental imagery that is facilitated by the same psychological mechanisms that usually subserve the production and reception of real speech. Although this Ryle-inspired theory involves the redeployment of existing psychological faculties to subserve belief self-ascription, it does not involve the redeployment of one’s reasoning faculties in the way imagined by the transparency theorist.

There is furthermore room to worry that the economy of the redeployment theorist’s explanation has been oversold. If psychological mechanisms deployed in ordinary inference are redeployed for the purposes of belief self-ascription, these mechanisms must somehow be insulated from the broader psychological profile that accompanies ordinary inferences. First, for reasons that we will see in Section 4, the disposition to deploy these mechanisms in the service of belief self-ascription must be insulated from dispositions to reason hypothetically and to increase one’s confidence in the conclusion of an inference continuously with increases in confidence in the premise. Second, when the psychological mechanisms for belief self-ascription are deployed, they must be prevented from leading to the standing psychological state of believing that one believes that *p* on the basis that *p*. Suppose for illustration that I believe on Monday that it will rain on Tuesday, and that this belief leads me to believe both that I believe on Monday that it will rain on Tuesday and also that the tennis match will be cancelled on Tuesday. Since my inference from the premise that it will rain to the conclusion that the match will be cancelled leaves me in the standing mental state of believing the conclusion on the basis of the premise, if I later go on to revise my belief that it will rain, I also will revise

¹⁵ Ryle’s suggestion appears in (1949, pp. 181-185), but see (Carruthers, 2011) for a more recent development of the idea and (Byrne, 2012 and forthcoming b) for helpful commentary. See also (Gordon, 2007), who proposes an alternative explanation in which our faculties for producing speech are redeployed in a different way to subserve belief self-ascription.

my belief that the tennis match will be cancelled.¹⁶ I will not revise, however, my belief that I believed on Monday that it will rain on Tuesday, because the psychological process that led me to the self-ascription, whatever it was, does not typically leave one in the standing mental state of believing that one believes that p on the basis that p. Without a better understanding of the psychology of ordinary reasoning than anyone currently has, it remains a highly speculative matter how economically this can be accomplished.

A final common source of motivation for Transparency comes from *considering examples*. While I think this source of motivation runs into problems, there are different kinds of problems with different kinds of examples. Here we will confine ourselves to one especially influential example. In his discussion of self-knowledge, Gareth Evans remarked that

[I]n making a self ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outward—upon the world. If someone asks me 'Do you think there is going to be a third world war?', I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question 'Will there be a third world war?'.¹⁷

While Evans' description of the example is plausible, we should be cautious about drawing conclusions about the nature of self-knowledge from examples featuring the use of psychological vocabulary in everyday discourse. For everyday discourse often features questions ostensibly about your mental states whose principal aim is not really to elicit a report about your mental states. When a stranger asks me 'Do you know how to get to Washington Square Park?', I must attend, in answering her, to precisely the same outward phenomena I would attend to in answering 'How do you get to Washington Square Park?' But that is because the latter question is the one she really wanted answered. If she instead was simply curious about the state of my geographical knowledge, I could tell her that I know how to get to Washington Square Park without bothering to think about how to get there. In Evans' example, it is natural to interpret the friend as prompting Evans to contribute to a discussion about whether there will be a third world war, rather than as soliciting a yes or no answer about whether Evans has an opinion. For this reason, claims about the appropriate procedure for responding to the friend's question, however plausible, are a poor guide to the procedure for knowing one's beliefs.¹⁸

¹⁶ Of course, there are psychological limits on our capacity to keep track of the evidence on which our beliefs are based. This point is made forcefully by Harman (1986), although I think he overstates the severity of these limits.

¹⁷ (Evans, 1982, p. 225).

¹⁸ For a dissenting view, see (Gordon, 2007). Although Gordon draws inspiration from Evans, his theory has some important differences from those discussed here.

3. Transparency and Reliability

Despite its attractions, Transparency can seem intuitively puzzling. According to our ordinary understanding of what a reason for belief is, p can be a good reason or basis on which to believe that q only if p is strong evidence for the truth of q . This ordinary understanding seems incompatible with Transparency. Setting aside special cases where you are a hyper-expert about whether p , the fact that p is not itself very strong evidence that you believe that p .

Below I will present a series of problems for Transparency, which I take to clarify and reinforce this intuitive puzzle. But first I want to set aside a related worry to which the Transparency theorist may have an effective reply. Although the worry is loosely inspired by ones that have appeared in the literature,¹⁹ I will present it in my own way, with the aim of highlighting what I see as the broader theoretical issues at stake. Although the worry itself will be set aside, we will take away some lessons about one kind of theory about inferential justification which serves as a natural compliment to Transparency, and which we will have a chance to discuss in greater depth in Section 5.

We can start with some common intuitions about Gettier-style examples like the following:

Broken Barometer: The barometer is broken, but you believe, on the basis of extensive supporting evidence, that it is working properly. You can see that the barometer reads ‘low’, and you believe on this basis that the barometric pressure is low.

Cases like Broken Barometer tend to elicit two intuitive reactions. The first is that, since you have every reason to believe that the barometer is working, it is *rational* or reasonable for you to believe that that the pressure is low. Indeed, when you see that the barometer reads ‘low’, it would be irrational for you not to believe that the pressure is low. The second reaction is that, since the barometer is broken, you are in no position to *know* that the pressure is low based on what you have seen—even if what you believe happens to be true.

I am taking Transparency to be the claim that p is a good reason for one to believe that one believes p , in the sense that one can be *justified* in believing that one believes p based on p . The worry I want to set aside, which concerns the Transparency theorist’s ability to explain our *knowledge* of our beliefs, might be stated like this. Since our self-ascriptions of belief not only are justified but furthermore amount to knowledge, any theorist with a story about how our self-ascriptions are justified is committed to a corresponding story about how they amount to knowledge. The Transparency theorist, in particular, is committed to the claim that one can know that one believes p , and not merely be justified

¹⁹ See, e.g., (Byrne, 2005 Sec. 7.2). A related objection is raised by Silins (2012, pp. 304-305).

in believing that one does, on the basis that p. And the lesson of cases like Broken Barometer, so the objection goes, is that one can know that q on the basis that p only if p is an objectively reliable indicator that q. This is incompatible with the Transparency theorist's proposal, however, because in typical cases p is not an objectively reliable indicator that one believes p.

This objection makes a few potentially controversial moves, but I think there is one in particular that the Transparency theorist is well-positioned to respond to. Let's say that p is an **adequate ground** for you to know that q just in case you are in a position to know that q on the basis that p. A crucial step in the objection is left-to-right direction of a first pass view about inferential knowledge that says that, for all values of p and q,

Evidential Reliabilism: p is an adequate ground for you to know that q
iff p is an objectively reliable indicator that q.

The Evidential Reliabilist owes us a story about what it means for p to be an objectively reliable indicator that q. But to a first approximation, it has something to do with the counterfactual sensitivity or statistical correlation between the p-facts and the q-facts. The guiding idea is that in order for one to have knowledge grounded in evidence, there must be some sort of causal or statistical connection between the evidence and what that evidence enables one to know. For example, in the case of the Broken Barometer the Evidential Reliabilist says that one cannot know that the pressure is low on the grounds that the barometer reads 'low' because the fact that it reads 'low' does not reliably indicate that the pressure is low—since the barometer is broken, its readings are not connected to the facts in the way that is needed to confer knowledge.

In reply to the Evidential Reliabilist's objection, the Transparency theorist cannot plausibly maintain that in general the fact that p is an objectively reliable indicator that one believes that p. Unless you are a weather hyper-expert, the fact that it will rain is not an objectively reliable indicator that you believe that it will rain. So the Transparency theorist must reject Evidential Reliabilism. Nonetheless, the Transparency theorist has available a different account of the Broken Barometer case, which replaces Evidential Reliabilism with the claim that for all values of p and q,

Inferential Reliabilism: p is an adequate ground for you to know that q
iff inferring from p that q is a reliable belief-forming process.²⁰

In most cases, the only realistic way for the process of inferring from p that q to be reliable is for the facts about p to be appropriately connected with the facts about q. But in the kinds of cases the Transparency theorist is concerned with, the process of inferring from p that one believes p can achieve its reliability in an unusual way. Since one must

²⁰ See, e.g., (Alston, 1986 and 1988) and (Comesaña, 2010) for more refined versions of Inferential Reliabilism, but note that Comesaña identifies his view as a kind of evidentialism.

believe that p in order to infer from p that q, the process of inferring from p that one believes p will be reliable, regardless of the connection between the p-facts and the facts about one's beliefs.²¹

Not only does this move help to forestall an objection to Transparency, it furthermore may be seen as providing an additional line of motivation in its favor. For consider the right-to-left direction of Inferential Reliabilism, which says that you can know that q on the basis that p if inferring from p that q is a reliable belief-forming process. Assuming that knowledge requires justification, this means that you can be justified in believing that q on the basis of p if inferring from p that q is reliable, as it is in the case of inferring from p that you believe that p.

We should be clear that Inferential Reliabilism is only a first pass view, and that those who are attracted to what it says in broad outline will wish to refine it in various ways (e.g., by adding a 'no defeaters' clause). Still, I think it is plausible that to the extent that the Inferential Reliabilist wishes to distinguish his view from Evidential Reliabilism, we should expect further refinements to leave the implication of Transparency intact. This is because a transparent inference is a paradigmatic case where a transition from the belief that p to the belief that q is reliable even in the absence of a corresponding reliable connection between the p-facts and the q-facts. So to the extent that one is attracted to Inferential Reliabilism, one should take this attraction to provide an additional line of motivation for Transparency. (The flip side, of course, is that to the extent that Transparency suffers from problems, so does Inferential Reliabilism—a point we will return to in Section 5.)

4. Transparency and Reasons for Belief

We have just considered and set aside a worry concerning the suitability of p as grounds for knowing that one believes that p. Put roughly, the worry was that for p to be grounds for knowing q, there must be an *objective* connection of an appropriate kind between the facts about p and the facts about q, and that the facts about the world and the facts about your beliefs are not so connected. It is time to start pressing my main line of objection to Transparency, which concerns more directly the suitability of p as a reason on which one can justifiably believe that q. Again put roughly, the worry is that for p to be a good reason to believe that q, there must be an appropriate kind of *subjective* connection between p and q, such that you might reasonably take the fact that p to show that q is true. This is a problem for Transparency because prima facie it does not seem true that you might reasonably take the fact that it will rain to show that you believe it will rain.

In rough outline, this source of discomfort with Transparency is a common one. But we need to be careful in spelling it out in detail. Although others have rightly noted a problem here, I do not think they have developed the problem in its strongest form.

²¹ Byrne (2005) makes a similar point in his response to what he calls "the puzzle of transparency."

Before attempting to do so, I will first consider three existing attempts, which I call the **independence objection**, the **nonstandard form objection**, and the **reflective endorsement objection**.

Start with the independence objection raised by Richard Moran. Although Moran accepts a view closely related to Transparency, he cautions that

the claim of Transparency *is* something of a paradox: how can a question referring to a matter of empirical psychological fact about a particular person be legitimately answered without appeal to evidence about that person, but rather by appeal to a quite independent body of evidence?²²

While I am sympathetic to Moran's felt sense of paradox, I think what he says is inadequate as it stands. We are often justified in inferring a conclusion from a body of evidence that concerns an independent subject matter, as when we infer from a premise about the reading of a barometer to a conclusion about the barometric pressure in the atmosphere. Of course, it might be that in this case the evidence from which one draws the inference is not really independent of the conclusion in the way Moran has in mind. But characterizing the relevant kind of independence is not a trivial task. It might be suggested that the facts about the world are not statistically or counterfactually associated with the facts about one's beliefs in the way that the facts about a reliable barometer's readings are associated with the barometric pressure. But as we already have seen in Section 3, the Transparency theorist has an effective reply to anyone who would require the premise of an inference to be an objectively reliable indicator of its conclusion. Another suggestion is that in the case of the barometer you have additional 'connecting evidence' that supports that the facts about the barometer's readings are relevant to the barometric pressure, and that this evidence is lacking for an inference from *p* to 'I believe that *p*'. The problem with this suggestion is that it is implausible that one must always have such connecting evidence. Such a general requirement leads to an implausible skepticism about inferential justification. Take the total stock of evidence that you know by non-inferential means. If your beliefs about other a posteriori subject matters are justified, they must be justified by inference from this evidence. But by stipulation, you have no additional connecting evidence available. Note that this skeptical problem requires no controversial assumptions about what can be known non-inferentially. Under a traditional Cartesian model of perception, this will include only facts about your own mental states. And the resulting skepticism will extend to all of your beliefs about the external world. But even if we extend the domain of non-inferential knowledge beyond what the Cartesian allows—to include ordinary perceptual knowledge, for example—we still are left with an unacceptable skepticism about everything that you do not know directly by perception.

²² (Moran, 2003, pg. 413), as quoted in (Byrne, 2005)

Turn now to the nonstandard form objection raised by Andre Gallois.²³ Gallois' worry is that an inference from the premise *p* to the conclusion that one believes that *p* does not fit any standard form of good inference. It is not inductively valid, nor is it inductively strong, nor is it an instance of inference to the best explanation. But I think this objection to Transparency is also inadequate. For it relies on the assumption that a justified inference must fit one of these standard forms. And this assumption is rejected by many theorists who might otherwise be sympathetic to our rough line of objection to Transparency. Consider for example an 'anti-reductionist' view about testimony, which says that you can be justified in believing that *q* based on a source's testimony that *q* even though you have neither inductive nor deductive reasons to believe that if the source says that *q* then *q* is true.²⁴ Anti-reductionists typically deny that justified inferences must always fit standard deductive, inductive, or abductive forms. But perhaps the anti-reductionist could still admit of some further sense in which you must take the source's saying that *q* to show that *q* is true. A similar point could be made about theories of perceptual justification that hold that one can justifiably infer conclusions about the external world from premises about one's experiences, even when the inferences do not fit the standard forms.²⁵

Finally, consider the reflective endorsement objection raised by Matthew Boyle. Calling an inference from *p* to 'I believe that *p*' not only "deeply unreasonable" but "mad," Boyle explains that

To believe that I believe *P* is to hold it *true* that I believe *P*. Being a reflective person, I can ask myself what grounds I have for holding this true. The answer '*P*' is obviously irrelevant. I am asking what shows that the proposition *I believe P* is true, and a modicum of rational insight will inform me that, even if it is true that *P*, this by itself has no tendency to show that I believe it. ... A (person-level) inference is not a mere transition from a stimulus to a response; it is a transition of whose terms I am cognizant, and whose occurrence depends on my—in the normal case: persistently—taking there to be an intelligible relation between these terms. This is what makes it possible for an inference to leave me with a

²³ (Gallois, 1996, pg. 47)

²⁴ See (Adler, 2006) for a helpful survey on the epistemology of testimony. I discuss one kind of anti-reductionism about testimony in (Barnett, MS b).

²⁵ Note that such theories do not need to hold that we actually do regularly engage in such inferences. With this caveat in mind, I think most theories of perception that I elsewhere call 'Cartesian' (Barnett, MS ???), would qualify. Notable exceptions include Vogel (1990), who thinks that we can infer conclusions about the external world from premises about our own perceptual experiences via inference to the best explanation, and Descartes himself, who thinks that there is a deductive argument to such conclusions that appeals to supplementary premises that can be known by intuition.

sustained belief: I can reflect on why I draw a certain conclusion, and when I do, I can see (what looks to me to be) a reason for it.²⁶

Boyle's criticism of Transparency appeals to two key ideas. The first, which I accept, is that you are justified in inferring from *p* that *q* only if you take the fact that *p* to show that *q* is true. The second is a kind of reflective endorsement requirement on justified inference, which holds that you are justified in inferring from *p* that *q* only if on reflection you would judge that *p* is a good reason to believe that *q*.²⁷ As I understand Boyle, his strategy is to appeal to the second of these ideas in support of the first. I have my doubts about whether the second idea really does support the first, and I argue elsewhere that the two ideas are moreover incompatible with each other.²⁸ We do not need to wade into this issue here, however. For there is a more direct way of seeing that no plausible reflective endorsement requirement poses a serious problem for Transparency. Note that any such requirement must help itself to a fair bit of idealization. Consider, for example, the student who is led by some hair-brained epistemological theory to believe that she has no good reason to believe that she has hands. On pain of counting the belief that she has hands unjustified, the reflective endorsement requirement must say that what matters is not the student's actual reflective endorsement, but instead her reflective endorsement under an idealization where she does not believe the hair-brained epistemological theory. With this in mind, it should be clear that the reflective endorsement requirement poses no special problem for Transparency. For the Transparency theorist thinks that *p* is a good reason to believe that one believes *p*, and rejects that this requires *p* to amount to evidence that one believes *p*. So the Transparency theorist can say that under idealized circumstances, an agent will be in a position to reflectively endorse *p* as a good reason to believe that she believes *p*, for such an agent will not mistakenly believe that this requires *p* to be evidence that she believes that *p*. (Of course, I agree with Boyle that these claims from the Transparency theorist should be resisted, but the point at issue here is whether the reflective endorsement requirement gives us any additional ammunition for doing so.)

So although I am sympathetic to the rough and intuitive idea behind these objections to Transparency, I think we need to exercise greater caution in developing the objection in greater detail.²⁹ Below, I will outline three closely related objections to transparency that all grow out of the guiding idea that you cannot justifiably believe a conclusion based on a premise unless you can justifiably take the premise to show that the conclusion is true.

²⁶ (Boyle, 2011, pp. 230-231)

²⁷ Indeed, Boyle wants to make the stronger claim that one cannot hold a belief at all, much less a justified one, without meeting this requirement. But we need not saddle the objection to Transparency with this stronger and more controversial claim.

²⁸ See (Barnett, MS ???) for further discussion.

²⁹ Other critics have stressed a more specific objection to Transparency stemming from a principle I call 'Inheritance', which I discuss below in Section 4.2.

After going through the objections individually, we will return in Section 5 to the bigger picture.

4.1 Hypothetical Reasoning and Indicative Conditionals

Our first problem for Transparency concerns its incompatibility with an appealing claim about the relationship between hypothetical and categorical reasoning, which says that inferential support is preserved under hypothetical supposition. Let me explain.

Consider a typical case of categorical reasoning: You see that the barometer reads ‘low’, and you infer that the pressure is low. In making this inference, you in some sense take it to be a fact that the barometer reads ‘low’, and you furthermore take this fact about the barometer's reading to show that the barometric pressure is low. Because inferring from p that q requires taking it to be a fact that p , the Transparency theorist can plausibly claim that you cannot infer from p that q unless you believe that p . This allows the Transparency theorist to cite the reliability of inferring from p that you believe that p in an effort to explain how your self-ascriptions of belief might amount to knowledge.

Not all reasoning involves making inferences in this sense, however. For even if you do not believe that the barometer reads ‘low’, you still can consider the possibility that it reads ‘low’ as a hypothetical supposition. If it is reasonable or rational for you to take the fact that the barometer reads ‘low’ to show that the pressure is low, then it should be rational for you to accept that the pressure is low under this hypothetical supposition. More generally, it is appealing to say that

Hypothetical Reasoning: If p is a good reason for you to believe that q , then you are in a position to accept³⁰ that q under the hypothetical supposition that p .³¹

If we accept Transparency, however, we must reject Hypothetical Reasoning. For if you consider the possibility that the number of stars is even as a hypothetical supposition, you could not accept under this supposition that you believe that the number of stars is even. While I take this claim to be *prima facie* plausible, it is worth reinforcing in greater detail. We can do so by considering a few important applications of hypothetical reasoning.

Our first application is the Ramsey Test for the assertability of indicative conditionals such as ‘If the barometer reads ‘low’, then the pressure is low’. According to a famous suggestion from Frank Ramsey,

³⁰ You are in a position to accept something if it rationally permissible for you to do so. Being in a position to accept something does not require the psychological ability to do it.

³¹ See (Cohen, 2010) for a recent defense of Hypothetical Reasoning, although Cohen's main concerns are quite different from those at issue here.

If two people are arguing ‘if p will q?’ and both are in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q.³²

One natural way of developing Ramsey’s suggestion, which replaces his heuristic talk about two people arguing with talk about a single person reasoning, says that

Ramsey Test: ‘If p, then q’ is assertable for you iff you can accept that q under the hypothetical supposition that p.

We have on the table three claims. The first claim is Transparency, which concerns the justification of self-ascriptions of belief. The second claim is Hypothetical Reasoning, which concerns the relationship between hypothetical and categorical reasoning. And the third claim is the Ramsey Test, which concerns the relationship between hypothetical reasoning and the assertability of indicative conditionals. Put these three claims together, and we are left with absurd results. For if we accept both Transparency and Hypothetical Reasoning, then it follows from the right-to-left direction of the Ramsey Test that you can assert the conditional ‘If p, then I believe that p’ for any proposition p. But that is absurd. For instance, it would mean that the following conditionals are assertable: ‘If there is life on other planets, then I believe that there is life on other planets’, ‘If the number of stars is even, then I believe that the number stars is even’. Accepting every conditional of this form would be tantamount to accepting that you are an omniscient god. Adapting a slogan from David Chalmers and Alan Hájek, we can say that Ramsey + Transparency = God.³³

It would be a mistake to see this as a problem for the Ramsey Test, rather than for Transparency. For other familiar applications of hypothetical reasoning will lead us to equally absurd results.

Consider Reasoning by Cases. You know that either there is life on other planets or there isn’t. If you can accept that you believe there is life on other planets under the supposition that there is, then Reasoning by Cases will allow you to conclude ‘Either I believe that there is life on other planets, or there is not life on other planets.’ In short, Reasoning by Cases + Transparency = God.

The same goes for *reductio ad absurdum*. Suppose for *reductio* $\langle p$, but I do not believe that $p \rangle$. Under this supposition, Transparency will allow you to accept that you believe that p and you don’t believe that p. Since this is a contradiction, you can conclude that it is not the case that p and you do not believe that p. For example, you can conclude that it

³² (Ramsey, 1931)

³³ Invoking Moore’s Paradox, Chalmers and Hájek’s (2006) slogan is that Ramsey + Moore = God.

is not the case that there is life on other planets and you do not believe that there is. Reductio + Transparency = God.³⁴

The takeaway lesson is that if we accept Transparency, we must reject Hypothetical Reasoning. This is a significant cost,³⁵ although it is a cost that at least some Transparency theorists may be prepared to accept.³⁶

4.2 Continuity and Degrees of Confidence

My second objection to Transparency concerns its implications for self-ascriptions of belief under conditions of first-order uncertainty—such as in a case where your meteorological evidence leaves it open whether it will rain, and you cannot justifiably hold a belief either way. Before stating the objection directly, I first wish to consider a pair of closely affiliated objections that have made appearances in the recent literature. The first of these objections, as discussed by Michael Martin, Sydney Shoemaker, and Nico Silins, appeals to the idea that when you believe that *q* on the basis that *p*, your belief that *q* will *inherit* any epistemic defects possessed by your belief that *p*.³⁷ For example, if you believe that the tennis match will be cancelled on the basis that it will rain, then your belief that the match will be cancelled will be unjustified if your belief that it will rain is unjustified. Generalizing from cases like this might lead us to accept the more general principle that

Inheritance: If you believe that *q* on the basis that *p*, then your belief that *q* is unjustified if your belief that *p* is unjustified.

To see why Inheritance is in tension with Transparency, bear in mind that if Transparency is to help us explain the justification of ordinary self-ascriptions of belief, then ordinarily we must believe that we believe that *p* on the basis of *p* itself. According to Inheritance, this means that in ordinary cases your belief that you believe that *p* will be justified only if your belief that *p* is justified. Yet it seems that when you self-ascribe beliefs in the

³⁴ The incompatibility of Transparency with *reductio ad absurdum* reasoning has been noted independently by Valaris (2011).

³⁵ Hypothetical Reasoning has been questioned recently by Weatherson (2012), but for reasons far removed from those at issue here. Weatherson claims that when reasoning non-deductively we can arrive at absurd conclusions using only hypothetical reasoning and other plausible rules of reasoning. But as Dogramaci (2010) argues, the reasoning that allows us to reach the absurd conclusion also involves the use of a Universal Introduction rule that is highly suspect in the context of non-deductive reasoning.

³⁶ See especially Byrne's (2005, Sec. 7.2) comparison with the rule of necessitation in modal logic, which allows you to move from the lemma that *p* to the conclusion that necessarily *p*, so long as the lemma that *p* was derived in an appropriate fashion. I think it is independently implausible that the rule of necessitation is a good rule of reasoning in the sense that is of interest here. (Of course, if you know that you have derived that *p* in an appropriate fashion, then *this* fact may be a good reason on whose basis to believe that necessarily *p*.)

³⁷ See (Martin, 1998), (Shoemaker, 2009) and (Silins, 2012, pp. 304-305).

usual way, your self-ascription can be justified even if the self-ascribed belief is unjustified.

Although I am sympathetic to this objection in spirit, I think that as it stands it faces two serious problems. The first is that insofar as Inheritance is supported solely by induction from uncontroversial cases, it is open to the Transparency theorist to reject Inheritance in favor of an alternative principle that gets the uncontroversial cases right without causing problems for Transparency.³⁸ The second problem is that the Transparency theorist might take issue with the claim that self-ascriptions of unjustified beliefs are themselves unambiguously justified, and claim instead that they are merely *locally* justified. Let me explain.

Suppose you know that the tennis match will be cancelled if it rains. If you form an unjustified belief that it will rain, there seems to be an important sense in which, given the fact that you believe this, you ought to believe that the tennis match will be cancelled. This is what I have in mind when I say that a belief is locally justified—rough speaking, the belief fits in with its immediate doxastic surroundings. Even so, intuitively there is a further *global* sense in which the belief that the tennis match will be cancelled is unjustified. What it takes for a belief to be globally unjustified is not obvious, but one initial gloss is that a globally unjustified belief is one that would be locally unjustified if the attitudes in its immediate doxastic surroundings were justified. The belief that the tennis match will be cancelled qualifies as globally unjustified in this sense because if you were to doubt as you should that it will rain, then the belief that the tennis match will be cancelled would be locally unjustified. If this initial gloss is on the right track, then the Transparency theorist could reply to the objection from Inheritance by claiming that, like your belief that the tennis match will be cancelled when you unjustifiably believe that it will rain, the self-ascriptions of an unjustified belief is locally justified but globally unjustified. For if you unjustifiably believe that it will rain, then there surely is a sense in which, given that attitude, you ought to believe that you believe that it will rain. But it also is true that if you withheld belief as you should from the proposition that it will rain, then you would not be even locally justified in believing that you believe that it will rain. Although I have my doubts about the gloss on global justification on which this reply depends, rather than contesting it I think we are better off pressing our objection to Transparency in a way that avoids cases of self-ascribing attitudes that are themselves unjustified.

Turn now to the second affiliated objection, which concerns cases where in response to inconclusive meteorological evidence you justifiably remain uncertain whether it will rain. In such cases it seems clear that you can both know and justifiably believe that you are uncertain whether it will rain. But the objection, advanced recently by Brie Gertler, is that the Transparency theorist seems unable to explain how your belief that you are

³⁸ Byrne (2005) develops a plausible reply of this sort to a closely related objection that concerns knowledge rather than justification.

uncertain is justified.³⁹ More generally, it seems that the Transparency theorist is able to explain only how you are justified in self-ascribing beliefs, and not how you are justified in self-ascribing other doxastic attitudes like uncertainty.

Although I am sympathetic to this objection as well, it is less forceful than we might wish because so long as our objection appeals solely to Transparency's lack of explanatory power, it remains open for the Transparency theorist to expand his theory to handle a wider range of cases. For example, it might be possible for the Transparency theorist to claim that the proposition that it is unclear whether *p* is a good reason to believe that you are uncertain whether *p*. It might even be possible simply to deny that the explanation of our self-ascriptions of doxastic attitudes is uniform, and claim that the justification of our self-ascriptions of uncertainty is different in kind from that of our self-ascriptions of belief.

Although I am skeptical that these replies on the Transparency theorist's behalf ultimately succeed, there is a more decisive problem stemming from examples involving first-order uncertainty. For the Transparency theorist seems not only unable to say the right thing about such examples, but moreover committed to saying the wrong thing. The difficulty stems from the rough thought that if *p* itself is a good reason for you to believe that *q*, then the strength of your epistemic position with respect to *q* should vary continuously with that of your position with respect to *p*. Roughly speaking, what this means is that the stronger your reasons are for believing that *p*, the stronger your reasons will be for believing that *q*. For example, if the fact that it will rain is itself a good reason for you to believe that the tennis match will be cancelled, then the stronger your reasons are for believing that it will rain, the stronger your reasons will be for believing that the match will be cancelled. Again put roughly, the claim is that

Continuity: If *p* is a good reason for you to believe that *q*, then the stronger your reasons are for believing that *p*, the stronger your reasons are for believing that *q*.

While I take the intuitive motivation for Continuity to be apparent, there are some important delicacies involved in spelling it out in greater detail.⁴⁰ For our purposes, we can restrict our attention to a particular consequence of Continuity: If *p* is a good reason for you to believe that *q*, then you can be at least roughly as confident that *q* as you are that *p*.

Let's first see how we get from the general claim of Continuity to this particular consequence. First, even if *p* is a good reason for you to believe that *q*, you can always have other reasons for believing that *q*. So your reasons for believing that *q* can be stronger than your reasons for believing that *p*. The second point of clarification

³⁹ (Gertler, 2011)

⁴⁰ For discussion, see (Kotzen, forthcoming), (Pryor, 2004, pp. 350-351), and (White, 2006, p. 532).

concerns the sense in which your reasons for believing that *p* can come in varying degrees of strength. While I think this notion should be understood in probabilistic terms, for our purposes here we can be quite modest in our commitments. The crucial point for our purposes is that if your reasons for believing that *q* are at least as strong as your reasons for believing that *p*, then you are in a position to be at least as confident that *q* as you are that *p*.⁴¹ For example, if you have strong enough evidence to be pretty sure that it will rain, then, assuming that this is a good reason for you to believe that the tennis match will be cancelled, Continuity will say that you are in a position to at least be pretty sure that the tennis match will be cancelled. Finally, since believing that *q* does not entail subjective certainty that *q*, it is plausible that *p* can be a good reason to believe that *q* even if certainty that *p* leaves room for a sliver of doubt as to the truth of *q*. Correspondingly, in cases where you are not certain that *p* it may be appropriate for you to be slightly less confident that *q* than you are that *p*.

The Transparency theorist is in no position to accept this consequence of Continuity. For the epistemic relationship between *p* and the proposition that you believe that *p* is not continuous in the way it requires. For example, if you have strong enough evidence to be pretty sure that it will rain, this does not mean that you can at least be pretty sure that you believe that it will rain. For if you are only pretty sure that it will rain, then presumably you are in a position to know that you do not believe that it will rain.

4.3 Alternative Elimination

The final problem for Transparency concerns its incompatibility with the claim that:

Alternative Elimination: *p* is a good reason for you to believe that *q* only if you have prior justification to reject the alternative possibility that *p* but not-*q*.

We can take rejecting (or eliminating) a possibility to consist in believing that it does not obtain. And we can take prior justification to mean justification that does not depend on your believing that *p* or having justification to believe that *p*. If so, it should be clear that the Transparency theorist must reject Alternative Elimination. For in cases where you do not believe that *p* and lack justification to believe that *p*, you often are in no position to reject the possibility that both *p* and you don't believe that *p*. For example, you are in no

⁴¹ A potential worry here, emphasized to me by an anonymous referee, is that one cannot always be as confident of the remote consequences of one's evidence as one is of the evidence itself because of higher-order doubts about one's own reliability at determining what the consequences of one's evidence are. But I think we can set this worry aside, since Continuity speaks only to cases where *p* is a good reason to believe that *q* in the (stipulative) sense that you are in a position to justifiably believe that *q* on the basis that *p*. Continuity is therefore silent about putative cases where *p* entails *q*, but where you are not in a position to justifiably believe that *q* on the basis that *p* due to higher-order doubts about your own ability to determine whether *p* entails *q*.

position to reject the possibility that the number of stars is even and you don't believe that it is even.⁴²

In rejecting Alternative Elimination, the Transparency theorist may appear to be in good company. A number of theorists, including some externalists about inferential justification, have rightly noted that they are committed to rejecting it. Still, I think a closer look at the kinds of moves available to such theorists cuts against this initial appearance of agreement. When externalists about justification contest Alternative Elimination, their concern is with its implications for cases like the following:

Mysterious Barometer: The barometer is working properly, but because you lack evidence one way or the other, you are uncertain whether it is working. You can see that the barometer reads 'low', and you believe on this basis that the barometric pressure is low.

This is the sort of case over which internalists and externalists about justification tend to disagree. The orthodox internalist verdict says that the fact that the barometer reads 'low' is not a good reason for you to believe that the barometric pressure is low. The orthodox externalist verdict says that it is a good reason for you to believe that the pressure is low, at least in the absence of various undermining conditions.

If we accept Alternative Elimination, then we must reject the orthodox externalist verdict. To see why, consider your epistemic situation prior to learning that the barometer reads 'low'. Since you are uncertain whether the barometer is working, you are in no position to reject the possibility that the barometer's reading, whatever it is, is inaccurate. In the absence of independent evidence that the pressure is low, therefore, you will be in no position to reject the more specific possibility that

Inaccurate 'Low': The barometer reads 'low', but the pressure isn't low.

This is why a theorist must reject Alternative Elimination if she wishes to say that the barometer's reading 'low' is a good reason for you to believe that the barometric pressure is low. In rejecting Alternative Elimination, such a theorist must choose between two unappealing options.

The first option involves denying the general principle that justification is closed under deductive implication. A first pass version of this general principle says that you are justified in believing a proposition only if you have justification to believe all of its logical consequences. This first pass version of the principle may need to be refined in

⁴² It is worth noting that Alternative Elimination plausibly follows from Hypothetical Reasoning using the *reductio ad absurdum* rule. Here I wish to offer a motivation for Alternative Elimination that is independent of Hypothetical Reasoning.

various ways,⁴³ but for our purposes we need only consider the restricted claim that for all values of p and q ,

Closure: You are justified in believing that q only if you have justification to reject that p and not- q .

The theorist who seeks to reject Alternative Elimination by rejecting Closure must say something funny about your epistemic situation after you learn that the barometer reads 'low'. For this theorist says that you are justified in believing that the pressure is low, and yet unjustified in rejecting the possibility that the barometer reads 'low' inaccurately.⁴⁴ Since it is obvious that the barometer cannot read 'low' inaccurately if the pressure really is low, it seems incoherent and unjustifiable for you to believe that the pressure is low even while remaining open to the possibility that the barometer reads 'low' inaccurately. So the theorist who denies Closure in ordinary situations like this seems committed to rampant epistemic quagmires, in which you cannot justifiably remain open to the possibility that the barometer reads 'low' inaccurately even though you cannot justifiably reject that possibility either.⁴⁵ This is a funny thing for a theorist to say. Nevertheless, saying this funny thing may be a cost some externalists are willing to pay.⁴⁶ Let's press on.

The second option involves denying that reasons for belief must help one to discriminate between what is true and what is false. In rough and impressionistic terms, the view that reasons must be discriminating in this way says that a reason to believe a proposition must take the form of evidence or considerations pointing in favor of its truth, while a reason to reject that proposition must take the form of evidence or considerations pointing the other way. If we accept this rough idea, then it seems that we also should accept the more specific claim that p can give us no reason at all to favor the possibility that p and q over the possibility that p and not- q —for regardless of which of these two possibilities obtain it will be the case that p . And if we accept the claim that p can give us no reason at all to favor the possibility that p and q over the possibility that p and not- q , then then we also should accept the more modest claim that p cannot give us sufficient

⁴³ I have in mind worries related to the failure of “logical omniscience,” which concern primarily the distant and unobvious consequences of propositions that we justifiably believe. For recent discussion, see (Christensen, 2004) and (Schechter, forthcoming).

⁴⁴ Even if a Closure-denying theorist says that you are not justified in rejecting that the barometer reads 'low' inaccurately, she still might grant that you are justified in rejecting that the pressure is not low. For comparison, a Closure-denier might say that you are unjustified in rejecting that you are a brain in a vat with no hands, but still allow that you are justified in rejecting that you have no hands.

⁴⁵ See the discussion of the Existence thesis in (Barnett, MS a).

⁴⁶ For a recent survey of attempts to reject Closure, see (Avnur, 2012). Although Avnur is not an externalist, he rejects Closure in special cases involving the rejection of extreme skeptical possibilities, allowing that in such cases you cannot justifiably remain open to possibilities that you also cannot justifiably reject.

reason to altogether reject the possibility that p and not-q. That is, for all values of p and q such that p and not-q is possible,⁴⁷

Discrimination: p is not a good reason to reject that p and not-q.

The theorist who seeks to reject Alternative Elimination by rejecting Discrimination says something funny about the epistemic impact of learning that the barometer reads 'low'. For this theorist says that when you look over to see that the barometer reads 'low', you thereby become justified in rejecting the possibility that it reads 'low' inaccurately. Since it is obvious that the barometer can read 'low' inaccurately only if it reads 'low', this too is a funny thing to say. For again, reading 'low' is exactly what you should expect from a barometer that reads 'low' inaccurately. So it is difficult to see how you could be justified in rejecting the possibility that the reads 'low' inaccurately on the basis that it reads 'low'.⁴⁸

We are now in a position to see the Transparency theorist's dilemma with respect to rejecting Alternative Elimination. For suppose that, in the face of limited or conflicting meteorological evidence, you are uncertain whether it will rain. Since you do not believe that it will rain, you are in no position to reject the possibility that

Unexpected Rain: It will rain, but you don't believe that it will rain.

Since you lack justification to reject the possibility that it will rain unexpectedly, it follows from Alternative Elimination that the fact that it will rain is not a good reason for you to believe that you believe it will rain. This does not mean that if you later come to believe that it will rain, you will be unable to justifiably self-ascribe the belief that it will rain. But it does mean that you will be unable to do so on the basis that it will rain. For this reason, the Transparency theorist must reject Alternative Elimination. And this is hard to do.

In contrast to the externalist about the Mysterious Barometer, the Transparency theorist cannot plausibly take the option that involves rejecting Closure. When the externalist denies Closure, the point is to make room for your being justified in believing that the pressure is low, even though you are not justified in rejecting the possibility that the barometer reads 'low' inaccurately. In contrast, it is not plausible that in ordinary cases of self-ascribing beliefs, you are justified in believing that you believe that it will rain and

⁴⁷ The requirement that p and not-q be possible is arguably necessary because it is arguable that p is a good reason to reject the proposition that p and not-p, as well as other propositions where the value of q is some other proposition incompatible with p. There is some room for disagreement over how best to understand the relevant notion of possibility. I am undecided between an understanding where the possibilities include only those things consistent with your evidence and one where they include everything that cannot be ruled out a priori. I do not think we need to decide the issue here. I thank an anonymous referee for discussion of this point.

⁴⁸ My understanding of this issue is indebted to White (2006), and to numerous discussions with Jim Pryor.

yet not justified in believing that it will not rain unexpectedly. So whatever we think about the plausibility of an externalist's rejecting Closure in order to avoid Alternative Elimination, a similar move is not available to the Transparency theorist.

It therefore seems necessary for the Transparency theorist to take the option of rejecting Discrimination, and of claiming that the fact that it will rain is itself a good reason for you to reject the possibility that it will rain unexpectedly. This move is not without its intuitive costs. In order for it to rain unexpectedly, it will have to rain. So it is difficult to see how you could be justified in rejecting the possibility that it will rain unexpectedly on the basis that it will rain.

More generally, the suggestion that one can be justified in rejecting the possibility that p and not- q on the basis that p is difficult to square with a natural conception of what a reason for belief is, which says that a reason to believe a proposition must take the form of evidence or considerations pointing in favor of its truth and conversely that a reason to reject the same proposition must take the form of evidence or considerations pointing the other way. There is much room for internal disagreement among theorists broadly sympathetic to this natural conception of epistemic reasons. But if we reject Discrimination, it is difficult to hold on to this natural conception of a reason for belief, even in broad outline. In order for it to be the case that p and not- q , it will have to be the case that p . So there appears to be no room for a view which says that p itself can point against the possibility that p and not- q .

4.4. Summary

It is time to take a step back, and consider the bigger picture. We began with a rough and intuitive objection to Transparency that appeals to a broadly evidentialist conception of epistemic reasons. According to the evidentialist, p can be a good reason for you to believe that q only if p is strong evidence for the truth of q . Since in typical cases p is not strong evidence that you believe that p , this means that p cannot be a good reason for you to believe that p . The more specific objections to Transparency considered in this section all grow out of this guiding idea.

It might therefore be worried that the preceding objections do not provide any *independent* motivation for rejecting Transparency, aside from the rough and intuitive idea that we began with. If the Transparency theorist is willing to reject the evidentialist conception of reasons from the outset, the worry goes, then it does no good to mount a series of objections that are by design closely aligned with this sort of evidentialism.⁴ But I think this worry is misguided, for at least two reasons. First, the objections can strengthen the case against Transparency by making its anti-evidentialist implications more concrete. It is easier for the Transparency theorist to reject a theoretical claim like evidentialism in the abstract than it is to reject evidentialism's most intuitively appealing

⁴⁹ I thank Alex Byrne and David Chalmers for making this point forcefully to me in conversation.

consequences or components. Second, there are reasons to be interested in the evidentialist objection to Transparency that go beyond our interest in persuading those who are not already convinced that Transparency is false. As I will now explain in Section 5, Transparency lies at the fault line between competing conceptions of inferential justification, one of which gives explanatory priority to inferential relations between one's beliefs, and the other of which gives priority to evidential relations between the objects of those beliefs. Even for those of us who find Transparency's anti-evidentialist commitments implausible in the abstract, a closer look at these commitments—and, in particular, their close relationship to unappealing consequences of orthodox reliabilism—can shed some light on this broader debate.

5. Reasons and Reliable Inferences

One important dimension along which theories of inferential justification can vary is the familiar one running between internalist theories and externalist ones like reliabilism. Another dimension of variation cutting across this concerns whether explanatory priority is given to evidential relations between facts or propositions, or to transitions between the mental states that represent those facts or propositions. If we are impressed by the preceding problems for Transparency, then I think we should take these problems also to favor our assigning explanatory priority to evidential relations rather than to transitions between mental states.

Recall that one motivation for Transparency stems from Inferential Reliabilism, which on a first pass says that p is an adequate ground for knowing that q just in case inferring from p that q is a reliable belief-forming process. As we saw in Section 3, Inferential Reliabilism stands in contrast to Evidential Reliabilism, which, instead of appealing to the reliability of the transition between the belief that p and the belief q , appeals to the reliability of the connection between the p -facts and the q -facts. Because the transition from the belief that p to the belief that you believe that p is an instance of a reliable process even when the fact that p is not a reliable indicator that you believe that p , these two views have conflicting implications for Transparency.

Corresponding to this distinction between Inferential and Evidential Reliabilism, we might draw a similar distinction between two kinds of internalist theories. While internalists have sought motivation for their view from a variety of quarters, one of its motivations comes from dissatisfaction with the verdicts that reliabilist views generate in cases like the Mysterious Barometer. As we have seen, the orthodox reliabilist verdict says that you can be justified in believing that the pressure is low based on the barometer's reading, even though you lack justification to believe that the barometer is reliable. Internalists who find this intuitively unsatisfactory often voice their complaint in rough terms as follows: In order for you to be justified in inferring that the pressure is low from the fact that the barometer reads 'low', there must be an appropriate kind of

subjective connection between p and q, such that you might reasonably take the fact that p to show that q is true.⁵⁰

While I think this complaint is on the right track, we must be careful in how we understand it. Consider two possibilities. On a first pass evidentialist construal, we can take this to mean that for all values of p and q,

Evidential Internalism: p is a good reason for you to believe that q iff you have justification to believe that p is an objectively reliable indicator that q.

On a first pass inferentialist construal, we might replace the requirement that one have justification to believe in an appropriate relationship between the p-facts and the q-facts with the requirement that one have justification to believe merely that the inferential transition from a belief that p to a belief that q is a reliable one. This gives us:

Inferential Internalism: p is a good reason for you to believe that q iff you have justification to believe that inferring from p that q is a reliable belief-forming process.⁵¹

The difference between inferentialist and evidentialist approaches concerns how to explain why some inferences are justified and others are not. The inferentialist prioritizes the reliability, or the subject's awareness of the reliability, of the psychological transitions between beliefs, while the evidentialist prioritizes evidential support relations between facts or propositions, and explains the justifiedness of transitions between beliefs in terms of them. The difference between these competing explanatory models is rarely noted, and when it is, it is sometimes taken for granted that inferential and explanatory models need not conflict about which inferences are justified.⁵² It is not hard to see why. As was the case with Evidential and Inferential Reliabilism, Evidential and Inferential Internalism give similar verdicts about most ordinary cases. For in ordinary cases, the transition from the belief that p to the the belief that q is an instance of a reliable process only if there is a reliable connection between the p-facts and the q-facts. And this means

⁵⁰ See, e.g., (BonJour, 1985, Ch. 3) and (Boghossian, 1989, pp. 8-9).

⁵¹ As with Evidential and Inferential reliabilism, Evidential and Inferential internalism must be further qualified to be plausible, for example by including a 'no defeaters' clause. With this qualification in place, I think it is plausible that Evidential Internalism, but not Inferential Internalism, has Hypothetical Reasoning, Continuity, and Alternative Elimination as consequences. However, Evidential Internalism is arguably stronger than these three theses in ways that might bring unwanted consequences. Although I am generally sympathetic to Evidential Internalism, I worry in particular about unwanted and unnecessary skeptical consequences. For this and other reasons, I regard both Evidential Internalism as merely a first pass view, and I hope to develop it further in future work. I think the qualifications necessary to avoid implausible skeptical consequences are unlikely to affect the present discussion.

⁵² See, e.g., (Alston, 1986).

that one's having justification to believe in the reliability of the transition ordinarily goes along with one's having justification to believe in the reliability of the connection.

However, there is once again an exception in the case of the transition from the belief that *p* to the belief that you believe that *p*. Arguably but plausibly, anyone with the requisite concepts, time, and attention is in a position justifiably to believe that the process of inferring from *p* that one believes that *p* is a reliable one.⁵³ Yet it is only in special cases of hyper-expertise about *p* that one has justification to believe that the fact that *p* is a reliable indicator that one believes that *p*. So in most cases, the Inferential Internalist will license as justified an inference from *p* to the conclusion that one believes that *p*, while the Evidential Internalist will not.

Two problems therefore confront the Inferential Internalist. The first is that Inferential Internalism's right-to-left direction has Transparency as a consequence. If the objections presented above lead us to reject Transparency, then we must reject Inferential Internalism as well. Now it could be replied on the inferentialist's behalf that Inferential Internalism, as I have characterized it, is only a first pass view, and that those who are attracted to what it says in broad outline will wish to refine it. Even if so, there is reason to worry that any refinement that avoids having Transparency as a consequence will leave in place an important explanatory problem. Evidential Internalism gives us an elegant explanation of why inferences from the premise that *p* to the conclusion that one believes that *p* are typically unjustified: It is because one typically lacks justification to believe that *p* is a reliable indicator that one believes that *p*. Even if Inferential Internalism is modified to avoid licensing such inferences as justified, there is reason to worry that any modification sufficient to do so will be objectionably ad hoc. Inferences from *p* to the conclusion that one believes that *p* are a paradigmatic example of a kind of inference where the verdicts of Inferential Internalism and Evidential Internalism come apart. Any modification to Inferential Internalism that is tailored to make it agree with Evidential Internalism on these paradigmatic cases seems bound to be ad hoc.

The second problem for Inferential Internalism concerns the motivation for its left-to-right direction, which gives as a necessary condition for one to be justified in inferring from *p* to *q* that one have justification to believe that this inference is a reliable one. To see the problem, we will need to consider briefly what the traditional motivation for this necessary condition is. Traditionally, internalists have been struck by the intuitive illegitimacy of certain kinds of inferences that are licensed by orthodox reliabilists. For example, concerning Mysterious Barometer it strikes many internalists as intuitively illegitimate for you to infer that the pressure is low from the fact that the barometer reads 'low'. If you lack evidence concerning the reliability of a barometer, these internalists say, then the mere fact that the barometer is objectively reliable will not mean that you

⁵³ For further discussion, see (Shoemaker, 1996, esp. Chs. 2 and 11) and (Setiya, MS). Note that the claim is only that anyone with the relevant capacities is in a position to have a justified belief that this inferential process is reliable, not that most people have this justified belief.

are justified in making such an inference. The necessary condition imposed by Inferential Internalism is intended to rule this inference unjustified—and it plausibly does so, since without justification to believe that the barometer is reliable, it seems you lack justification to believe that this inference is an instance of a reliable process. The problem for the Inferential Internalist again concerns the explanation of why this inference is unjustified. Intuitively, the reason why you cannot justifiably infer that the pressure is low from the barometer’s reading is that you cannot justifiably take the premise that the barometer reads ‘low’ to show that the conclusion that pressure is low is true. And Inferential Internalism fails to capture this intuitive explanation. This is demonstrated by the fact that it does not similarly rule as unjustified an inference from p to the conclusion that one believes that p , even though this inference is afflicted by the very same intuitive problem. For as we have seen, even though the transition from a belief that p to the conclusion that one believes that p is a reliable one, it does not confer justification because you cannot justifiably take the fact that p to support that you believe that p . Since the problem with this inference is the very same as the one afflicting the inference from the barometer’s reading to the conclusion that the pressure is low, any principle that adequately captures the problem with the latter inference ought to rule the former unjustified as well. But Inferential Internalism does not do this. For as we have already seen, the inference from p to the conclusion that you believe p meets the necessary condition imposed by Inferential Internalism. The upshot is that even though Inferential Internalism correctly rules unjustified the inference from the barometer’s reading to the conclusion that the pressure is low, it does not properly explain why this inference is unjustified. It is not unjustified because you fail to have justification to believe that the inference is reliable, as the Inferential Internalist says. Instead, it is unjustified for the same reason that typical inferences from a proposition p to the conclusion that you believe that p are unjustified—because you are not justified in taking the premise of the inference to show that the conclusion is true.⁵⁴

⁵⁴ I wish to thank Eli Alshanetsky, Dorit Bar-On, Ned Block, Paul Boghossian, Alex Byrne, Fabrizio Cariani, David Chalmers, Justin Clarke-Doane, Sinan Dogramaci, Dan Fogal, Kit Fine, Don Garrett, Sandy Goldberg, Tim Maudlin, John Morrison, Kate Nolfi, Elliot Paul, Jim Pryor, Stephen Schiffer, Nishi Shah, Nico Silins, Declan Smithies, David Velleman, Jared Warren, and audiences at the Syracuse Philosophy Annual Workshop and Network (SPAWN), Northwestern University, and the NYU dissertation seminar.

References

- Adler, Jonathan (2006) 'Epistemological Problems of Testimony' *The Stanford Encyclopedia of Philosophy*, Edward Zalta ed., available at <http://plato.stanford.edu/archives/spr2013/entries/testimony-episprob/>
- Alston, William (1986) 'Internalism and Externalism in Epistemology' *Philosophical Topics* 14(1): 179-221.
- Alston, William (1988) 'An Internalist Externalism' *Synthese* 74: 265-283.
- Avnur, Yuval (2012) 'Closure Reconsidered' *Philosophers' Imprint* 12(9).
- Aydede, Murat (2002) 'Is Introspection Inferential?' in *Privileged Access: Philosophical Accounts of Self-Knowledge*, Brie Gertler ed., Ashgate Publishing.
- Bennett, Jonathan (2003) *A Philosophical Guide to Conditionals*. Oxford: Oxford University Press.
- Boghossian, Paul (1989) 'Content and Self-Knowledge' *Philosophical Topics* 17(1): 5-26. BonJour, Laurence (1985) *The Structure of Empirical Knowledge*. Harvard University Press. Boyle, Matthew (2009) 'Two Kinds of Self-Knowledge' *Philosophy and Phenomenological Research* 78(1): 133-164.
- Boyle, Matthew (2011) 'Transparent Self-Knowledge' *Supplementary Proceedings of the Aristotelian Society* 85(1): 223-241.
- Brueckner, Anthony (1998) 'Moore Inferences' *Philosophical Quarterly* 48: 366-369.
- Byrne, Alex (2005) 'Introspection' *Philosophical Topics* 33: 79-104.
- Byrne, Alex (2010) 'Recollection, Perception, Imagination' *Philosophical Studies* 148: 15-26. Byrne, Alex (2011) 'Transparency, Belief, Intention' *Supplementary Proceedings of the Aristotelian Society* 85(1): 201-221.
- Byrne, Alex (2012) 'Review of Carruthers, *The Opacity of Mind*, Notre Dame *Philosophical Reviews*.
- Byrne, Alex (forthcoming a) 'Knowing What I Want' in *Consciousness and the Self: New Essays*, J. Liu and J. Perry eds, Cambridge University Press.
- Byrne, Alex (forthcoming b) 'Knowing that I am Thinking' in *Self-Knowledge*, Anthony Hatzimoysis ed., Oxford University Press.
- Carruthers, Peter (2011) *The Opacity of Mind: An Integrative Theory of Self-Knowledge*, Oxford. Cassam, Quassim (2011) 'Knowing What You Believe' *Proceedings of the Aristotelian Society* 111(1): 1-23.
- Chalmers, David and Hájek, Alan (2006) 'Ramsey + Moore = God' *Analysis* 67: 170-172. Christensen, David (2004) *Putting Logic in its Place: Formal Constraints on Rational Belief*, Oxford University Press.
- Christensen, David (2007) 'Epistemic Self-Respect' *Proceedings of the Aristotelian Society* 107(1): 319-337.
- Coliva, A. (2008) 'Peacocke's Self-Knowledge' *Ratio* 21(1): 13-27.
- Cohen, Stewart (2010) 'Bootstrapping, Defeasible Reasoning, and *A Priori* Knowledge' *Philosophical Perspectives* 24: 141-159.
- Comesaña, Juan (2006) 'A Well-Founded Solution to the Generality Problem' *Philosophical Studies* 129: 27-47.

- Comesaña, Juan (2010) 'Evidentialist Reliabilism' *Noûs* 44(4): 571-600.
- Dogramaci, Sinan (2010) 'Knowledge of Validity' *Noûs* 44: 403-432.
- Dretske, Fred (1994) 'Introspection' *Proceedings of the Aristotelian Society* 94" 263-278.
- Dretske, Fred (1995) *Naturalizing the Mind*. MIT Press.
- Edgley, Roy (1969) *Reason in Theory and Practice*. London: Hutchison & Co.
- Evans, Gareth (1982) *The Varieties of Reference*. Oxford: Oxford University Press.
- Fernández, Jordi (2003) 'Privileged Access Naturalized' *Philosophical Quarterly* 53(212): 352-372.
- Fernández, Jordi (2005a) 'Self-Knowledge, Rationality, and Moore's Paradox' *Philosophy and Phenomenological Research* 71(3): 533-556.
- Fernández, Jordi (2005b) 'Privileged Access Revisited' *Philosophical Quarterly* 55(218): 102-105. Gertler, Brie (2011) 'Self-Knowledge and the Transparency of Belief' in *Self-Knowledge*, Anthony Hatzimoysis ed., Oxford University Press.
- Gibbons, John (2006) 'Access Externalism' *Mind* 115(457): 19-39.
- Gordon, Robert (2007) 'Ascent Routines for Propositional Attitudes' *Synthese* 159: 151-165. Green, Mitchell and Williams, John (2007) *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*. Oxford: Oxford University Press.
- Gallois, Andre (1996) *The World Without, the Mind Within: An Essay on First-Person Authority*. Cambridge University Press.
- Harman, Gilbert (1986) *Change in View*, MIT Press.
- Kotzen, Matthew (forthcoming) 'Dragging and Confirming' *Philosophical Review*.
- Lycan, William (1999) 'Dretske on the Mind's Awareness of Itself' *Philosophical Studies* 95:125-133.
- Martin, Michael (1998) 'An Eye Directed Outward' in *Knowing Our Own Minds*, Crispin Wright, Barry Smith, and Cynthia Macdonald eds., Oxford: Oxford University Press.
- McHugh, Conor (2010) 'Self-Knowledge and the KK Principle' *Synthese* 173: 231-257.
- Moran, Richard (2001) *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.
- Moran, Richard (2003) 'Responses to O'Brien and Shoemaker' *European Journal of Philosophy* 11: 402-419.
- Peacocke, Christopher (1998) 'Conscious Attitudes, Attention, and Self-Knowledge' in *Knowing Our Own Minds*, Crispin Wright, Barry Smith, and Cynthia Macdonald eds., Oxford: Oxford University Press.
- Pryor, James (2004) 'What's Wrong With Moore's Argument?' *Philosophical Issues* 14: 349-378. Pryor, James (2001) 'Highlights of Recent Epistemology' *British Journal for the Philosophy of Science* 52: 95-124.
- Pryor, James (2005) 'There Is Immediate Justification' in *Contemporary Debates in Epistemology*, Matthias Steup and Ernest Sosa eds., Blackwell.
- Ramsey, Frank (1931) 'General Propositions and Causality' in *The Foundations of Mathematics and Other Logical Essays*, R. B. Braithwaite ed.
- Schechter, Joshua (forthcoming) 'Rational Self-Doubt and the Failure of Closure' *Philosophical Studies*.

- Schwitzgebel, Eric (forthcoming) 'Knowing Your Own Beliefs' *Canadian Journal of Philosophy*.
 Setiya, Kieran (2011) 'Knowledge of Intention' in *Essays on Anscombe's Intention*. Anton Ford, Jennifer Hornsby, and Frederick Stoutland, eds., Cambridge, MA: Harvard University Press.
- Shah, Nishi and Velleman, J. David (2005) 'Doxastic Deliberation' *Philosophical Review* 114(4): 497-534.
- Shoemaker, Sydney (1996) *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Shoemaker, Sydney (2009) 'Self-Intimation and Second-Order Belief' *Erkenntnis* 71(1): 35-51.
 Silins, Nicholas (2012) 'Judgment as a Guide to Belief' in *Introspection and Consciousness*, Declan Smithies and Daniel Stoljar eds., New York: Oxford University Press.
- Smithies, Declan (forthcoming) 'Moore's Paradox and the Accessibility of Justification' *Philosophy and Phenomenological Research*.
- Smithies, Declan (2012) 'A Simple Theory of Introspection' in *Introspection and Consciousness*, Declan Smithies and Daniel Stoljar eds., New York: Oxford University Press.
- Stoljar, Daniel (2004) 'The Argument from Diaphanousness' in *New Essays in the Philosophy of Language and Mind*, M. Ezcurdia, R. Stainton, and C. Viger eds., Calgary: University of Calgary Press.
- Tye, Michael (2002) 'Representationalism and the Transparency of Experience' *Noûs* 36(1): 137-151.
- Vahid, Hamid (2005) 'Moore's Paradox and Evans's Principle: A Reply to Williams' *Analysis* 65(4): 337-341.
- Valaris, Markos (2011) 'Transparency as Inference: A Reply to Alex Byrne' *Proceedings of the Aristotelian Society* 111(2): 319-324.
- Vogel, Jonathan (1990) 'Cartesian Skepticism and Inference to the Best Explanation' *Journal of Philosophy* 87(11): 658-666.
- Way, Jonathan (2007) 'Self-Knowledge and the Limits of Transparency' *Analysis* 67(3): 223-230.
- Weatherson, Brian (2012) 'Induction and Supposition' *The Reasoner* 6(5): 78-80.
- White, Roger (2006) 'Problems for Dogmatism' *Philosophical Studies* 131: 525-557.
- Williams, John (2004) 'Moore's Paradoxes, Evans's Principle and Self-Knowledge' *Analysis* 64(284): 348-353.
- Williams, John (2006) 'In Defense of an Argument for Evans's Principle: A Rejoinder to Vahid' *Analysis* 66(2): 162-170.
- Williams, John (2007) 'Moore's Paradox, Evans's Principle, and Iterated Beliefs' in *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*, Mitchell Green and John Williams eds., Oxford: Oxford University Press.
- Williams, John (forthcoming) 'Moore-Paradoxical Belief, Conscious Belief, and the Epistemic Ramsey Test' *Synthese*.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zimmerman, Aaron (2004) 'Unnatural Access' *Philosophical Quarterly* 54(216):

- 435-438. Zimmerman, Aaron (2006) 'Basic Self-Knowledge: Answering Peacocke's Criticisms of Constitutivism' *Philosophical Studies* 128(2): 337-379.
- Zimmerman, Aaron (2008) 'Self-Knowledge: Rationalism vs. Empiricism' *Philosophy Compass* 3(2): 325-352.