
DEFERRING TO OTHERS ABOUT ONE'S OWN MIND

BY

CASEY DOYLE

Abstract: Pessimists about moral testimony hold that there is something suboptimal about forming moral beliefs by deferring to another. This paper motivates an analogous claim about self-knowledge of the reason-responsive attitudes. When it comes to your own mind, it seems important to know things “from the inside”, in the first-personal way, rather than putting your trust in another. After motivating pessimism, the paper offers an explanation of its truth. First-person knowledge is distinctive because it involves knowing a state of mind and finding it intelligible from one’s point of view. It concludes by considering the value of this form of self-understanding.

1. Pessimism

According to “pessimists”, there is something objectionable or second best about forming moral beliefs by testimony.¹ As Allison Hills puts it, ‘[c]-oncerning a distinctively moral question ... it seems important to make up your own mind rather than put your trust in others’ (2013, p. 552). Pessimism is restricted to cases of ‘pure testimony’, that is, cases in which a subject unreflectively bases her belief on another’s word alone (McGrath, 2009). There is nothing problematic about listening to another, reflecting oneself, and coming to one’s own conclusion on the matter. Things are quite different, the thought goes, when you simply take another at her word. The problem is not, as some have claimed, that moral knowledge is not available by testimony.² Pessimists do not deny that one can gain knowledge if one defers to another on moral matters nor do they claim that one ought never defer (Howell, 2014). Rather, they claim that moral knowledge by testimony

is suboptimal. There is something we want when we want moral knowledge that cannot be transmitted by testimony. Furthermore, the problem with deference is not simply an instance of the relative demerits of testimony in general (Howell, 2014). There is something about moral knowledge in particular, which renders deference problematic.

I am a pessimist about self-knowledge. When it comes to your own mind, and in particular your own reason-responsive attitudes like belief, intention, desire, and the emotions, it seems important to know things from the inside or in the first-personal way, rather than putting your trust in others.³ The first-personal way is a method for knowing one's attitudes that is available only to the subject herself. Knowledge arrived at by this method is "first-person knowledge"; it contrasts with knowledge of one's mind arrived at by broadly third-personal means such as observation, inference from behavior, and testimony. I will not make any assumptions here about what the first-person method is, only that there is one. At times I will use the label 'introspection' for this sort of knowledge. But this is merely a label: I do not assume that self-knowledge is achieved by inner sense.

Pessimism about self-knowledge (hereafter "pessimism") can be captured as follows:

Pessimism: Other things being equal, it is preferable to know one's reason-responsive attitudes in the first-personal way rather than testimony.

It helps to consider examples. I will mention two, one involving the emotions and another involving belief.

Holiday Stress: Aaron and Bill are friends discussing their plans for the winter break. After Aaron details the itinerary of his vacation to the Caribbean with evident joy and excitement, Bill describes his own plans to spend time with in-laws in some dreary part of the country. As he mentions the uncertainties and frustrations associated with the trip, it becomes clear to Aaron that Bill is angry. Perhaps he is angry at his in-laws, at his wife, or at all of them. He is clearly angry at someone, and Aaron tells him this. Aaron is pretty reliable about this sort of thing. Without any further reflection, Bill accepts Aaron's testimony and forms the belief that he is angry.

Teacher Meeting: Jane is an expert about early child education and learning disabilities. Her own child is struggling in school. The school calls home and proposes to assess the child, as they believe she may have a learning disability. Does Jane believe that her child has a learning disability? Her partner, who knows her well, believes that she does believe this and tells her. Jane is not so sure. She knows what

she wants to believe but has a tough time keeping her imagination, anxiety, and desires in check when considering the matter. Jane unreflectively takes her partner's word for it that, deep down, she has this belief.

In both of these cases, the subject is ignorant of their mind and is, for whatever reason, prevented from gaining first-person knowledge of the relevant attitude. Often we can know our reason-responsive attitudes spontaneously, without any deliberation, "just like that". But, of course, not always. Pretty much everyone should grant this.⁴ We know from empirical psychology that we often make mistakes about our minds. And it is a perfectly familiar experience for a close friend to know what you think on some matter or how you feel better than you. In such a circumstance, it would seem to make sense to defer to another. You can form beliefs about your mind this way, and when all goes well, these beliefs count as knowledge. Indeed, there may be cases where, all things considered, you ought to defer to another. Even granting all this, intuitively, if we reflect on cases, we find that deference is second best. This means two things. First, there is something *distinctive* about first-person knowledge, a property it possesses that cannot be transmitted by testimony. Second, there is something *valuable* about this distinctive feature.

As should be clear, pessimists do not deny that others are ever in a better position to know your mind than you are. Far from it: Pessimism is supported by the theory and practice of talk therapy. People often go to therapy because they cannot make sense of themselves. And they go to a professional whose expertise consists, in part, in the ability to know and understand another's mind, to come up with an interpretation. The subject wants self-knowledge, and the therapist can possess the knowledge she seeks. You might expect that the therapist transmits this knowledge to the patient. But that is precisely not how it works. Here is Freud (1966):

From what I have so far said a neurosis would seem to be the result of a kind of ignorance – a not knowing about mental events that one ought to know of ... Now it would as a rule be very easy for a doctor experienced in analysis to guess what mental impulses had remained unconscious in a particular patient. So it ought not to be very difficult, either, for him to restore the patient by communicating his knowledge to him and so remedying his ignorance ...

If only that was how things happened! ... Knowledge is not always the same as knowledge: there are different sorts of knowledge, which are far from equivalent psychologically ... The doctor's knowledge is not the same as the patient's and cannot produce the same effects. If the doctor transfers his knowledge to the patient as a piece of information, it has no result ... The patient knows after this what he did not know before – the sense of his symptom; yet he knows it just as little as he did. Thus we learn that there is more than one kind of ignorance (pp. 347–348).

One of the aims of therapy is self-knowledge. And part of the therapist's expertise consists in the ability to know the patient's mind. But what

the patient wants in aiming at self-knowledge cannot be transmitted by testimony.⁵

I proceed as follows. The next section spells out pessimism in more detail. Section 3 argues that what is epistemically distinctive about self-knowledge cannot explain pessimism. Then I argue that the first-person perspective is distinctive because it involves, as I put it, occupying the perspective of the attitude known. After objecting to one way of understanding this, I present my own view. This leads to an explanation of the truth of pessimism and a consideration of why first-person knowledge is valuable to us.

2. *Motivating pessimism*

Pessimism is not restricted to testimony from another *person*, as is clear from the following remarks by Gaylin (2000):

More often than not an angry patient is defending himself against his own feelings of guilt – better to blame someone else than assume personal responsibility. This is a particularly crucial insight ... Were all of this knowledge available to a patient through the use of some elaborate diagnostic software package installed in one's home computer, I seriously question whether it could change fixed patterns of neurotic behavior ... Such abstract knowledge is likely to produce intellectualization, an enemy, not ally, of insight (p. 203).

Gaylin uses the nebulous but useful term 'insight' to refer to the form of self-knowledge that is the aim of therapy and that is especially valuable to us. This form of self-knowledge is not restricted to exotic mental states, and the idea of it is not hostage to abstruse psychoanalytic theorizing. Pessimism is part of our ordinary understanding of the first person, although it is especially clear in the therapeutic context. In this spirit, Moran (2001) locates a

crucial therapeutic difference between the merely 'intellectual' acceptance of an interpretation, which will normally be seen as a form of resistance, and the process of working-through that leads to a fully internalized acknowledgment of some attitude which makes a felt difference to the rest of the analysand's mental life. ... [W]hat is to be restored to the person is not just knowledge of the facts about oneself, but self-knowledge ... (pp. 89–90).⁶

Pessimism in other domains is typically motivated by eliciting an intuition that people have when they reflect on particular cases. Pessimists hope others share the intuition and then seek to explain it. In our case, examples of simply taking another at her word about whether one is angry or what one intends to do are meant to elicit that intuition. There is something odd about deference here that is absent in other cases. As with other forms of pessimism, not everyone will share the intuition. Some are *optimists*: they do not see what the big deal is with testimony. Others are *skeptics*: they do

not think knowledge of one's attitudes by testimony is so much as possible. What can we say to them?

I will start with the optimists. First, as we saw, we need not rely only on intuitions about cases because pessimism embodies a tenet of much talk therapy. Second, some will accept the explanans even if they are unsure about the explanandum. We can live with that. Third, we can note that pessimism is fairly weak. Pessimism holds that knowing your mind by testimony is suboptimal or second best. There is something we often want when we want self-knowledge that cannot be transmitted by testimony. If you are after that, there is reason for you not to simply defer. Furthermore, as mentioned, pessimism is only a claim about self-knowledge of the reason-responsive attitudes. Nobody should deny that there are many facts about one's mental life that are best known by testimony. The most obvious example is character traits like honesty or modesty. Pessimism is restricted to those features of one's mental life that are within the scope of first-person knowledge (which character traits are not) and in particular the attitudes.

What about skeptics? I suppose that if one thought that we are infallible or omniscient about some nontrivial class of our attitudes, then one would think that testimony about one's mind would never be necessary. But it seems safe to assume that view is false. Everybody should grant that there are circumstances in which another is better placed to know what one believes, desires, or what emotions one is feeling. We can be in the dark about our attitudes, even if this is only because our capacity to know them has been temporarily interfered with.

Anscombe (2000) argued that it is a necessary feature of intentional actions that the agent performing one possesses 'practical' knowledge of what she is doing. If you ask someone why she is ϕ -ing and she replies, 'Oh, I didn't realize that I was doing that!' then, according to Anscombe, other things being equal, it just follows that she was not ϕ -ing intentionally. This might suggest that it just does not make sense for someone to be in a position to need testimony about her own intentional actions. And one might extend the general picture to the attitudes. If someone needs to be told what she is doing, then she is not doing it intentionally; if someone needs to be told what she believes, then she does not really believe that. I do not doubt that there is an intimate connection, in creatures like us, between our capacity to hold certain attitudes and our capacity to know them (Shoemaker, 1994), but even as a claim about intentional action, this seems overly strong. There are surely lapses in one's awareness of what one is doing. And even if one finds this strong claim compelling in the case of intentional action, it is even more controversial as a claim about our reason-responsive attitudes.

As far as I can tell, there is no compelling case that knowledge by testimony about one's mind is not available. It seems safe to assume that

deferring to another, if they are trustworthy, would lead to knowledge. Our point is only that such knowledge is suboptimal.

One final worry. You might think that what is odd about deferring to another about one's mind is not deference *per se* but rather being in a position where one must rely on it. That is, it is strange to be in a position where you cannot just introspect to discover your attitudes. Our intuitions about cases like those described earlier are responsive to this strangeness, the sense that things are amiss with you, rather than the relative merits of introspection over testimony.⁷ This position is not skeptical, because it allows that knowledge of one's mind is available by testimony, and it endorses the intuition about cases like Bill and Jane. But it denies that these intuitions are sensitive to our preference for first-person knowledge.

This worry might be understood one of two ways. It might involve a normative claim about what sorts of thing one ought to know or it might involve a statistical claim about what we typically do know or are in a position to know. Compare: It is odd or second rate to form a belief about when your child was born by testimony. Why is that? Well, it might be because information about your child is something for which you are epistemically responsible: It is the sort of thing you should know. Or it might be odd simply because people typically know when their children were born. Neither claim has anything to do with testimony. It is odd or suboptimal to *need* testimony not to *gain* it.

Start with the normative reading. In the case involving your child, we expect you to have a bit of knowledge. But I do not know that we expect you to have that knowledge by any particular source. Suppose that for whatever reason you never noticed what day it was, and so, subsequently, you had to be told when your child was born. (Maybe the child was born right before or right after midnight.) A bit odd, perhaps, but it seems perfectly kosher from the perspective of satisfying your epistemic responsibilities: you know what you are supposed to know. Perhaps we expect people to know their minds; perhaps what we believe, desire, intend, and so on are the sorts of thing for which we are typically epistemically responsible. But the intuition behind pessimism is that we expect people to know their minds a particular way, the first-person way. Some evidence: the sense that things are second rate does not go away if we imagine a case in which someone already possesses self-knowledge by testimony, as opposed to forming a self-ascription by testimony. We do not expect people to know their minds, one way or another, but to know them from the inside. So the mere idea that we are responsible for knowing our minds will not capture the intuition behind our cases.

On the other hand, the statistical claim is either dubious or too weak to explain our intuitions. The claim is that, in the ordinary run of things, you are in a position to know your attitudes in the first-person way. In the normal case, it is perfectly obvious to you what attitudes you hold. This would mark an important difference between the cases of moral knowledge

and self-knowledge. It could not be seriously suggested that the oddity of deferring to another about, say, whether eating meat is permissible is explained by the fact that it is so very obvious what the answer is. The thought, then, is that the oddity of our cases is explained entirely by the fact that they are abnormal. It has nothing to do with an intuitive preference for one form of knowledge over another.

To put it bluntly: I deny that it is abnormal to be in a position where you are unsure about what you believe, what you want, and so on, and so where you might defer to another. So I do not think the intuition that something is funky with Bill and Jane can be chalked up to the thought that these are weird or *recherché* cases.

Notice that this is consistent with the claim that something is, in some sense, amiss with the subjects in these cases. For example, Moran (2001) claims that so long as one is rational, one will be in a position to gain first-person knowledge of one's attitudes. One might reasonably have doubts about this claim, especially when it comes to our emotions and desires, but leave them to one side. It is a familiar fact that we are very often less than perfectly rational and so very often in a position where, by Moran's lights, first-person knowledge is not available. The point is just that, on the one hand, it is doubtful that being in a position where one might rely on another's word is all that uncommon, and on the other, if being in such a position means that things have gone awry with one, this could only be in a way that, for creatures like us, they very often are.

I will assume from here on out that there is something intuitively suboptimal about deference and something distinctive and valuable about first-person knowledge. Looking forward, the three authors quoted above gesture at the distinctive feature of first-person knowledge that explains pessimism. As Freud says, self-knowledge is a distinctive form of knowledge that is different in its psychological features. Rather than focusing on methods, we should explain pessimism by identifying distinctive features of self-knowledge as a state of mind. As Moran puts it, what matters here is something like an 'internalized acknowledgment' that makes a 'felt difference'. And as Gaylin suggests, this acknowledgment is realized in understanding why one thinks or feels as one does. First-person knowledge is distinctive because it involves knowing an attitude and finding its possession to some degree intelligible from one's current point of view. This feature of self-knowledge cannot be transmitted by testimony, which is why we are pessimistic. But before spelling out this account, I will object to a natural alternative.

3. *The method and the status*

Philosophers writing about self-knowledge often mention two ideas: a distinctive method and a special status. There is a uniquely first-personal way

of knowing one's own mind, a way unavailable to anyone else. This is the method. Self-ascriptions arrived at by this method possess a special epistemic status. This can be understood in different ways. Perhaps the method is highly reliable, or the beliefs formed by it are highly justified and very safe, or they have infallible justification when all goes well. For our purposes, it will not matter exactly what the status is.

We cannot explain pessimism by appeal to the method or the status. Start with the method. The bare fact that I alone have access to a particular method cannot explain a preference for using that method – maybe the method has nothing to be said for it (Davidson, 1984). Suppose that your paranoid and senile great-grandmother is in a nursing home. She has lost some precious item and is convinced that one of the staff stole it from her. The staff want to know what has happened to the item, but your great-grandmother will not breathe a word to anyone there about her suspicions. She will only speak to you. So you alone have access to this method of forming beliefs about what happened to the item. But there is nothing to be said for that method.⁸

This might incline you to think that it is the status that matters. This seems plausible enough, at least in the ordinary run of things. Because you have access to a better method, you had better use it. Suppose, though, that you have a friend, Inez. She is infallible, but only about two things: your current mental states and the ambient temperature. I want to make two claims about this case. First, while there is nothing wrong with deferring to Inez about the temperature, there is something suboptimal about deferring to her about your beliefs, even though she is infallible. We might say: Inez is no substitute for introspection, therapy, or reflection. Second, we cannot appeal to the status to explain this. Even those who think that introspection provides factive reasons for belief deny that our self-ascriptions are, quite generally, infallible. But Inez is infallible. Suppose too that she is honest and that you know that she is infallible and honest. She is infallible about your mind and an impeccable testifier. Plausibly, from the point of view of the status, whatever it is, few methods are better than deferring to Inez, including the first-person one. Nevertheless, there is something suboptimal about deferring to her.⁹ Ruling out error is not the way to go if we want to explain our intuitive preference for first-person knowledge.

4. Consciously occupying the perspective of a mental state

Familiar ideas from the literature on self-knowledge cannot explain pessimism. Reliable diagnostic software or a therapist can only provide the patient with a bit of information about himself, or what Freud calls 'the sense of his symptom'. Access to that information by peculiar means or an especially secure grasp of it are no better off. What is needed is

something beyond information about one's own mind. The proposal to be advanced here is that it is a particular kind of *feeling* or *experience*, that of finding the attitude appropriate, intelligible, or presumptively authoritative. I think that this is the best way to capture a widely acknowledged phenomenon, namely that the first-person perspective involves a normative perspective on the attitude known. In this section, I set out the phenomenon and object to one way of understanding it.

Suppose that you are at a diner, notice a slice of pie enshrined in a rotating display case, and think 'I want some pie!' This thought expresses both self-knowledge and one's attraction to the pie. In having the thought about one's desire, one also, thereby, takes up a perspective on the object of the desire: that it is desirable. From your point of view, you cannot separate the awareness of the desire from the feeling of attraction to the pie. The question of the desirability of the pie is hard to separate from the question about whether you desire some. On a plausible account of the reason-responsive attitudes of creatures like us, they are commitments, stances on the world, or answers to questions (Hieronymi, 2009; Gibbons, 2013).¹⁰ In holding a reason-responsive attitude, a subject takes a stance on how the world is or how she will behave, and she is thereby responsible for defending that stance in the face of scrutiny. For example, to believe that p is to be committed to its being true that p ; to intend to ϕ is to be committed to ϕ -ing; and to desire to ϕ is to be committed to its being in some way good to ϕ . At a first pass, the phenomenon in question is that when one takes the first-person perspective on a reason-responsive attitude, one cannot be indifferent to the commitment it embodies. Taking the first-person stance on the attitude seems to involve taking the stance of the attitude on the world, or, as I will put it, it involves consciously occupying the perspective of the attitude known.

Many in the self-knowledge have a finger on this. Gallois (1999) argues that we can only possess first-person knowledge of beliefs that are subjectively justified. Moran (2001) claims that 'the expression of one's belief carries a commitment to its truth' (p. 92). Fernandez (2013) makes the observation that 'we feel pressured to occupy the mental states that we attribute to ourselves' (p. 14). This feeling is a phenomenological feature of self-ascriptions that he calls 'assertiveness'. Leite (2018) points out that in order for first-person knowledge to play a role in reflectively revising our attitudes, it must involve grasping them 'as our own'. Putting the idea in a handy slogan, Burge (1996) writes that, 'the first- and second-order perspectives are the *same* point of view' (p. 110, emphasis in original).¹¹

Of course, not all of these philosophers understand the phenomenon the same way. But they seem to agree that the first-person perspective involves consciously occupying the perspective of the attitude known. As I understand it, this is a claim about the phenomenology of self-knowledge. It captures what it is like to know one's mind in the first-personal way. This is strongly suggested by examples like the one involving you and that slice of

pie. It can also be highlighted by contrast. My son has all sorts of interesting beliefs about dragons, and I am well versed in his beliefs. But knowing what he believes does not involve taking up that perspective on dragons to any degree. Knowing that he believes that dragons live in caves leaves entirely open to me the question whether they, in fact, do. Of course, this holds equally when I do not have a view on the matter. My colleague tells me he believes that *p*. I have never considered the question whether *p*. I know that he believes that *p*. Knowing this, the question whether *p* is, from my point of view, entirely open. It might be the case that the fact that he believes that *p* provides me with a *prima facie* reason to believe that *p*. It might be that I have a bit of evidence for believing *p*. But merely knowing his mind does not involve taking up its perspective.

We should like to know how precisely to understand the idea of consciously occupying the perspective of a state of mind. By far the dominant view in the literature appeals to metaphors of agreement, endorsement, ratification, and unity (Heal, 1994; Burge, 1996; Peacocke, 1998; Moran, 2001; Boyle, 2011; Coliva, 2012).¹² I will call this the *reflective endorsement model*. On the model, in knowing one's mental states, one also executes a higher-order assessment of them and endorses them.¹³ First-person knowledge that one believes that *p* itself involves undertaking a commitment to the truth that *p*.¹⁴ In the terminology of answering questions, the first-order and second-order beliefs give the same answer to the same question, one directly and the other indirectly (Gibbons, 2013). At the heart of this view is the claim that the normative aspect of self-knowledge is captured by or simply realized in a normative judgment. One might accept this for different reasons. Some philosophers think that normative judgment is part of the process of arriving at self-ascriptions (Moran, 2001); others hold that our capacity for self-knowledge just is the capacity to form attitudes on the basis of our assessment of the reasons for them (Rödl, 2007; Boyle, 2011); others hold that 'belief is a matter of viewing a proposition as what one ought to believe' (Marcus, 2016, p. 375); and others think acts of self-ascription undertake or ratify first-order commitments (Heal, 1994; Coliva, 2012). For our purposes, these differences will not matter.

However it is arrived at, the model is committed to the claim that the first-person perspective is restricted to attitudes that conform to the normative judgment one would make about them. The problem with the model is straightforward: first-person access is not so restricted. The model predicts that it is impossible to possess first-person knowledge of attitudes one does not reflectively endorse. But it is common to know, in the first-personal way, desires one would never act on, or emotions that seem ungrounded, or beliefs that are resistant to deliberation.

Consider a case of recalcitrant emotion (Bar-On, 2004). While far from phobic, I am uneasy on planes. I often feel fear as a plane takes off. I realize

that this is not appropriate. I am aware that flying is far safer than many things I do on a daily basis without worry or regret. I do not endorse the fear and wish I did not experience it. One might be inclined to describe such a state as an alien presence in my mind. But while there might be cases like that, recalcitrance is not restricted to them. My own fear makes sense to me to a degree, and the perspective on the flight it represents is one that I occupy. I feel compelled to embrace the idea that flying is dangerous, despite what I know. I occupy this perspective while judging that it is fine to fly. In a word: it is a mess inside my mind, but I have a first-person perspective on this mess. Simply having an eye on the reasons would miss this, because it would exclude competing voices.¹⁵

5. *Self-understanding and feeling*

We should reject the reflective endorsement model because it presents us with a Manichean conception of the first-person perspective and cannot make sense of taking that perspective on attitudes that do not conform to one's current perspective on reasons. The first-person perspective constitutively involves a normative perspective on the attitude known. It is natural to take this to involve a normative judgment, such as a reflective judgment that one ought to hold the attitude or that one has reason to. But we can occupy the first-person perspective even if we would refrain from such a judgment. An alternative is suggested by a view from the literature on moral testimony.

On a prominent view, the problem with moral testimony is that, quite generally, testimony cannot transmit understanding (Hills, 2009, 2010). Testimony can transmit knowledge that *p*, but it cannot transmit understanding why *p*. According to Hills (2010), understanding why consists in the possession of a set of abilities: 'To understand why *p* you must have an ability to draw conclusions about similar cases, and to work out when a different conclusion would hold if the reasons why *p* were no longer the case' (p. 194). On this view, understanding must be distinguished from knowledge. One can know why something happened, know that *p* because *q*, without understanding it. That is because one can know the reason without grasping it as a reason.¹⁶ It might be thought that this offers a plausible account of what is distinctive of the first-person perspective, and so a plausible explanation of pessimism. On the proposal, what is distinctive about the first-person perspective is that it involves self-understanding, in particular understanding why one holds a particular attitude. Understanding cannot be transmitted by testimony, which is why we are pessimistic. This proposal seems especially well suited to explain Freud's observation. The patient does not just want to know her mind but to understand it.

Because understanding cannot be transmitted by testimony, the therapist cannot help the patient simply by offering testimony.

I think there is something right about this proposal, but it is only plausible if we focus on a particular form of self-understanding. Not all forms of understanding why a person holds an attitude involve commitment to the attitude as, in some way, appropriate. Again, I might understand why my son believes that dragons live in caves, or why my colleague believes that *p*, without thereby being committed or taking a stance on the relevant questions. Likewise, the form of understanding manifest in offering a rational explanation of an attitude is not first-personal; it is not something that can only be achieved by the subject herself. A natural amendment: the first-person point of view involves understanding an attitude by appeal to one's own reasons for holding it. In Moran's (2001) terms, when one knows that one believes that *p* in the first-personal way, one sees that belief as according with one's 'sense of reasons' for so believing (see also Parrott, 2017). The thought here is that when you know your attitude in the first-personal way you understand it, in the sense that would be manifest in offering one's own reasons for holding that attitude. It is because one has that explanation at one's disposal that one is able to occupy the perspective of the attitude known.

The ability to explain one's attitudes by appeal to one's own reasons is something only the subject can possess. Because it is a form of understanding, it cannot be transmitted by testimony. Furthermore, it is not the same thing as reflective endorsement. I can offer my reasons for something I have done without, here and now, endorsing them. Thus, this proposal looks well suited to explain pessimism. However, it is highly implausible to suppose that, quite generally, occupying the first-person perspective involves being able to offer a rational explanation of the attitude known. Two examples already mentioned demonstrate this. One might know that one desires the pie or that one fears flying without grasping one's reasons for doing so. Indeed, often part of the problem with recalcitrant attitudes is that one does not know why one holds them.

We can retain the basic idea here without over-intellectualizing the first-person perspective. From the first-person perspective one's attitudes make sense to one; they are intelligible in light of one's 'sense of the reasons'. But the idea of a sense of reasons need not be cashed out in terms of an articulate knowledge of the considerations on the basis of which one holds the attitude. Instead, the limit case for such understanding might be a particular *feeling*, the feeling of the attitude as making sense or as being compelling. Plausibly, this feeling is found when the attitude is manifest in conscious experience. Consider, again, the pie. Why is it difficult, from your point of view, to separate the awareness of the attitude from endorsement of the commitment it embodies? Arguably, it is because conscious desires are *presumptively self-authenticating*. When one's desire is manifest in

consciousness, one feels oneself compelled to go get the pie. This might be because desires present their objects as good or to be done (Raz, 2010). Likewise, in the ordinary case, conscious fear itself involves feeling that its object is to be avoided.

In the same spirit, consider the case of conscious judgment, or the manifestation of belief in conscious experience. It is plausible to think of it as a *feeling of rational compulsion with respect to a proposition*. Judging that *p* in consciousness involves finding oneself compelled to endorse that *p*. Although he focuses on belief, Travis (2012) captures the phenomenon well in the following passage:

To believe is to be under the sway, or influence, of something: (one's encountering of) things being as they are. To be under such influence is to *feel* it – to feel compelled to think as one thus does ... such felt compulsion must be of a special kind ... It must not be *psychological* compulsion that one sees himself as under (p. 396).

Travis describes this feeling as 'Lutheran': in judging that *p*, one feels one's cognitive hands tied, as it were, able to do no other (*ibid.*). When you have an experience like this, it strikes you that *p* is the thing to think.

We are trying to identify a form of self-understanding of one's attitude that does not require knowledge of one's reasons. My claim is that certain self-conscious experiences are like that. Attitudes like belief, desire, and emotion are sometimes manifest in conscious experience. These experiences are self-conscious in that one is aware of them in having them. For our purposes, we can leave it an open question what form this self-consciousness takes. (Our question is not how we know or become aware of our attitudes, but what such awareness is like such that pessimism is true.) When you have an experience like this, the attitude is presumptively self-authenticating; it presents itself to you as appropriate or *prima facie* warranted. Furthermore, these experiences have the normative perspective we are interested in: one is aware of the attitude known, occupies it, and takes it as appropriate to some degree. Yet finding it appropriate in this way does not require knowing the reasons on the basis of which one holds it. One can find one's conscious fear intelligible and, as we say, feel its force, without knowing why one holds it, and so without being able to offer a rationalizing explanation of it. Finally, occupying the perspective of the attitude in this way is not the same thing as reflectively endorsing it. That is because this form of self-consciousness is achieved merely by the conscious manifestation of the attitude, when nothing intervenes. Such an occurrence does not require reflective judgment about reasons. And you can have this experience, and so occupy the first-person perspective on the attitude, in the face of a reflective judgment to the contrary.

Let us recap. The first-person perspective is distinctive because it involves consciously occupying the perspective of the attitude known.

‘Consciously occupying’ cannot be reflectively endorsing because there are cases of the former where the latter is conspicuously absent. The idea of understanding is helpful, but consciously occupying the attitude cannot require understanding why in the form of an ability to offer a rationalizing explanation, either. Instead, I have claimed that we can occupy the perspective of an attitude when the attitude is self-consciously manifest in our experience. This is a limit case on self-understanding. When we are lucid, suitably attentive, and so on, we can be aware of our attitudes when they are manifest in consciousness. This involves finding the stance the attitude embodies on the world presumptively self-authenticating. That it is only presumptively so matters, because we can have a feeling like this without endorsing the stance that the attitude embodies, as happens in cases of recalcitrance.

Earlier, I mentioned that the appeal to understanding could make sense of Freud’s observation. Patients do not just want information about themselves, they want to make sense of themselves. It is not obvious, though, that what they need is an articulate grasp of their reasons. It is extremely plausible that at least part of what they aim for is a particular kind of *experience*, an experience in light of which the relevant self-ascription makes sense, even if they cannot (yet) articulate their reasons for holding the attitude. In this spirit, in an influential paper on insight in psychotherapy, Richfield (1954) discussed the importance of emotional insight into one’s condition, and he placed a special emphasis on self-conscious experience, writing: ‘insights which incorporate the actual, conscious experience of their referents can be termed “ostensive” insights. These are obtained through the direct cognitive relations involved in the acquisition of knowledge by acquaintance’ (p. 405, quoted in Wallerstein, 1995, p. 171).¹⁷

6. *Explaining pessimism*

Assuming that this plausibly characterizes the normative perspective that is built into the first-person perspective, two questions remain. First, why can’t this perspective be transmitted by testimony? Second, what is valuable about having this perspective? Answering the first question would explain why we might seek out first-person knowledge in particular, and the second would explain what we hope to gain from it.

Why can’t deferring to another lead one to consciously occupy the perspective of the mental state known? One answer is straightforward. I have said that consciously occupying the perspective of a mental state involves finding it intelligible and that this can occur through the self-conscious manifestation of the attitude. It is pretty clear, I take it, that deferring to

another does not lead to that. If I take another at her word that I am angry, I do not thereby, except by accident, self-consciously feel anger.

This is right, as far as it goes, but we can offer a more robust explanation. First-person knowledge involves knowing an attitude and finding it appropriate. Under what circumstances would the attitude of 'finding appropriate' be warranted? As we have seen, there are two. The first is when you have a justified belief about your reasons for holding the attitude, that is, when you are in a position to offer a rationalizing explanation of it. The second is when the feeling arises in the self-conscious experience of the attitude. Part of what it is to consciously desire to ϕ is to find it appropriate, to a degree, to do so. However, the fact that someone else tells you that you want to ϕ does not, of itself, warrant you in finding that desire appropriate. Stuart tells me he wants the world to stop. I learn something about his mind, but do not, thereby, find it in any way appropriate that the world stop. Testimony about someone's mind does not warrant one in taking that person's attitude as appropriate. But another's testimony about one's own mind is just testimony about someone's mind, someone who happens to be oneself. So if Stuart tells me that I want the world to stop, absent other reflection on the matter, it is not intelligible for me to take that attitude as appropriate, to any degree.¹⁸

7. *The value of self-knowledge*

In a host of circumstances, we want self-knowledge. And what we want cannot be transmitted by testimony, even by those who know us better than we know ourselves. Instead, we want something that can only be achieved by knowing our minds in the first-personal way, "from the inside". What we want must be some distinctive feature of first-person knowledge. Like other philosophers, I think that first-person knowledge is distinctive because it involves a normative perspective on the attitude known. Rather than reflective endorsement, I have spelled this out in terms of a feeling of finding the attitude appropriate or intelligible, a feeling that is found in the self-conscious experience of attitudes. The problem with testimony is that another's word is not a fitting basis for this feeling: being told that one holds an attitude is not sufficient for finding it appropriate. Making sense of ourselves is something we have to do ourselves.

Assuming all this is correct, we can still ask: why is first-person knowledge so understood valuable to us? I do not suppose there need be one reason. I am going to mention two, although the topic certainly merits more discussion.

The presumption in favor of first-person knowledge is felt by the subject herself, but it is also felt by others. If I want to know what you believe or what you want, I want you to speak from the first-person perspective. It just

will not do, in many circumstances, to hear it from someone else, or for you to pass it on second hand. That the first-person perspective involves a normative perspective can explain this quite neatly. Often, when I want to know what you believe, want, or feel, I do not just want to find out a fact about you, as when I want to know your blood type. Rather, I want you to take a stand on some first-order issue. For example, I want to know what you want because we are trying to figure out what to do today. I want to know what you believe because we are trying to come to a conclusion together or I am pitching an idea and I want to know what you think. I want you to engage with the first-order issues themselves and to speak your mind, as we say, rather than to simply report facts about your mind. Knowledge of one's mind based on testimony lacks this normative perspective; in reporting it, you do not take a stand on the first-order issues at all, just as you do not when you report on another's mind. And that is why, often, it just will not do.¹⁹

But of what value is this knowledge to the subject herself? An answer recommends itself if we notice that the reflective endorsement model is poorly posed to provide one. On the model, the first-person perspective involves endorsing the attitude known as justified, supported by reasons, or what one ought to hold. The first-person perspective would then be restricted to those attitudes that conform to one's present deliberations about reasons; everything else would be 'alien' and only accessible by third-personal means. Now, it is not clear why, from my own point of view, I need to know the attitudes that conform to my present judgments about my reasons. If I am rational, attending to new evidence automatically results in revisions to my attitudes without reflection. And if my attitudes conform to my current assessment of the reasons I have, why attend to the attitudes instead of the reasons? It might be true that, in these circumstances, I can answer the question what attitudes I hold by figuring out what attitudes I ought to hold, by my lights. But what would be the point? Here is a rough, but helpful, analogy. Suppose I am trying to rid my house of flies. There are two species in the house. One is such that, if a window is opened, it is guaranteed to make its way out. Another is not like that, it will sometimes mill about even when the windows are opened. It is clear that, if I am concerned about keeping my house in order, there is no need to go looking for the first kind of fly. Attitudes that conform to one's current normative judgments are like that. There is not much point in knowing them.

Of course, one might deny that first-person knowledge is of instrumental value. If first-person knowledge is restricted to attitudes that conform to one's judgments about reasons, then one will possess or be in a position to possess it so long as one is rational. Thus, one might say, self-knowledge itself is not good for the subject, but the state that makes self-knowledge possible, namely, rationality, is. This sits well with a central claim of

Moran's (2001) that the ability to gain first-person knowledge is a sign of psychic health.

I have two objections to the proposed explanation. First, it is not at all clear why conforming with my normative judgments, here and now, should be any indication of psychic health in an intuitive sense. What if my normative judgments are the product of self-deception, evasion, or confusion? (Lear, 2004). It is true that attitudes that conform lead to a more consistent overall state of mind, other things being equal. But it is not at all clear that psychic health and consistency are the same thing (Harcourt, 2017). Second, more can be said about the value of self-knowledge once we reject the reflective endorsement model, as I have argued we should.

We have been led to this thought: if first-person knowledge is of instrumental value to us, then it might be precisely in cases of attitudes that do not conform to our current normative judgments. I claimed earlier that it is a feature of such cases that they make sense to one to a degree, in virtue of their phenomenology, but that the subject typically lacks a rational explanation of them. One is not sure why one fears flying, given what one knows about the risks. It is natural to suppose that, at least sometimes, people want a fuller, more developed understanding of themselves. This suggests that achieving that kind of self-understanding for attitudes that do not conform to one's current normative judgments requires first-person knowledge. That is, it is only if we possess insight into our attitudes, understood as a conscious experience of finding them compelling, that we can develop a fuller understanding of them. First-person knowledge is valuable to us precisely because there is a value in stepping away from the perspective of our normative judgments and occupying the perspective of mental states that do not conform to the reasons we take ourselves to have. There may be multiple reasons for this. One, suggested by what was said earlier, is that our normative judgments might themselves be confused or the product of self-deception. This could be corrected by better understanding what seems stubborn and recalcitrant but might reflect what we really think and feel on the matter. Put simply: Perhaps it is better to put one's trust in the recalcitrant emotion or attitude than the reflective judgment. Wollheim (2003) nicely articulates the hazards of restricting oneself to one's current normative judgments:

The moral is this: if we try to change ourselves before we have come to recognize ourselves, learned to know ourselves, there is a clear danger. It is that we shall put one form of self-ignorance behind us only to embrace another. We shall become volunteers in the happy cause of hypocrisy (p. 35).

Keeping our mental houses in order requires more than reflection on reasons, because it is not obvious that we should achieve agreement between our attitudes by getting our attitudes to conform to our normative

judgments. Instead, at least sometimes, we might seek to understand our attitudes on their own terms. Psychic health requires something other than changing ourselves to achieve coherence; it requires that we recognize ourselves before we change ourselves. It requires understanding ourselves as we are. A full exploration of this is well beyond the scope of the present paper. What we have seen, and what is striking, is that this form of understanding cannot be transmitted by testimony, even by those in the know.

8. *Vindicating pessimism*

There are three stages of the pessimist project. The first is motivating the view by appeal to intuitions and, in our case, the practice of talk therapy. The second is explaining it by identifying a feature of the form of knowledge in question (moral, aesthetic, or self-knowledge) that cannot be transferred by testimony and which is, intuitively, important to us. The third is vindicating pessimism by revealing that this feature is, in fact, important, so much so that there should be a defeasible presumption against ways of knowing the relevant facts that lack it. An illustrative example: Mogensen (2017) argues that pessimism about moral testimony is explained by our concern for the value of authenticity. But Mogensen is ambivalent about the value of authenticity: it is just not clear that we should care all that much about it. I have sought to motivate and explain pessimism about self-knowledge. But I have not vindicated it. Of course, it might just be obvious that self-understanding as described here is important. But perhaps we, or some of us, overvalue it. (The insistence on the first-person way of knowing might come across as too strident, much like the insistence, from some, on the superiority of talk therapy to other ways of treating mental health issues.)

The question of whether pessimism can be vindicated is an ethical and empirical one. It is empirical because it depends, to some degree, on the extent of first-person knowledge, the obstacles that stand in its way, and our chances of overcoming them. Throughout the history of philosophy, many have been impressed by the Delphic injunction to 'know thyself'. But among those moved by the demand, some were hopeful about our chances of meeting it, while others have been quite skeptical.²⁰ The question is ethical because even if there is something to be said for first-person knowledge, it is not clear what place it has in a well-lived life. This paper has not sought to answer these questions. It has sought to show that making sense of ourselves in this way, whatever its importance, is something we must do ourselves.²¹

University of Oxford

NOTES

¹ Nickel (2001), Hopkins (2007), Hills (2009, 2010), Howell (2014), among others.

² See Hopkins (2007) for discussion.

³ It is controversial to claim that the emotions are reason-responsive, but I assume so in what follows.

⁴ Some philosophers deny that first-person knowledge amounts to a discovery of one's attitudes. Instead, on their view, part of what it is to hold certain attitudes is to believe that you hold them (Shoemaker, 1994; Rödl, 2007; Boyle, 2011). If this view is to be at all plausible – if it is not committed to denying the existence of familiar cases like those just canvassed – then it must hold that second-order beliefs can remain implicit, tacit, or masked at particular times. See Shoemaker 2012 for discussion.

⁵ You might be pessimistic about other sources of knowledge of one's mind, like observation or inference. However, I will restrict focus here to testimony for three reasons. First, doing so allows us to exploit parallels between self-knowledge and moral knowledge. Second, by doing so we avoid begging the question against observational and inferential models of self-knowledge. Third, the puzzle is especially apparent in the therapeutic setting, where testimony would be apt. Thanks to Kieran Setiya for raising this issue.

⁶ The passage specifies: self-knowledge 'that meets the transparency condition'. Because the idea of transparency will not be central to our discussion, I have not focused on it here.

⁷ Thanks to John Gibbons and an anonymous referee for raising this worry.

⁸ Presumably, self-knowledge satisfies a modal condition absent in the case of your great-grandmother. It is possible for another to form beliefs from her testimony (say, someone dressed up like you), but it is not possible for another to use the first-person method to form beliefs about your attitudes. Still, I think the example brings out the fact that having unique access to a method is not itself valuable to us. Thanks to Matthew Parrott for raising this wrinkle.

⁹ That self-knowledge has a special kind of epistemic status (e.g., immunity from error through misidentification, (Evans, 1982) or immunity from brute error, (Burge, 1996)) cannot explain pessimism, either, because, trivially, Inez is immune from the same errors.

¹⁰ I will use these phrases interchangeably. Like most philosophers who operate with the term, I do not have a definition of commitment to offer. The rough idea will suffice for our purposes.

¹¹ Of course, not everyone will accept the phenomenon, but I take it as a datum to be accounted for in a theory of self-knowledge. You might think that it mistakes the pragmatic function of self-ascriptions for a feature of the state of self-knowledge, for example (Martin, 1998).

¹² Not all of these authors endorse the reflective endorsement model for our attitudes across the board.

¹³ Some philosophers deny that self-knowledge involves a reflective higher-order belief that is caused by a first-order state (Shoemaker, 1994; Boyle, 2011). They think that the second order is constitutive of the first order. That is consistent with the reflective endorsement model as I understand it.

¹⁴ It need not be the case that this act of reflective endorsement is achieved by a stretch of conscious deliberation, as if one is aware of the belief, considers its credentials, and then endorses it. For example, one might hold that one possesses self-knowledge merely by or in holding an attitude and yet also think that self-knowledge involves reflective endorsement. Reflective endorsement is a way of characterizing the act of knowing one's mind, not necessarily a process of arriving at self-knowledge.

¹⁵ Leite (2018) raises a complaint in a similar spirit. If self-knowledge is sometimes involved in the revision of our attitudes, then it cannot require that we endorse the attitude known.

¹⁶ This is controversial, of course, but I will assume it for present purposes.

¹⁷ It is plausible to suppose, with Richfield, that the awareness achieved in the self-conscious manifestation of an attitude is something like Russell's knowledge by acquaintance. But we need

not assume that to agree with him about the importance of self-conscious experience. See Wallerstein (1995) for a helpful discussion of Richfield's views.

¹⁸ For a parallel treatment of Pessimism about aesthetic testimony, see Whiting (2015), and for moral testimony, see Doyle (forthcoming).

¹⁹ The distinction between an engaged or deliberative perspective on one's attitudes and a merely psychological one is a central theme in Moran (2001) and Roessler (2013).

²⁰ As Garrett (2017) points out, there is a long tradition in philosophy (including the Stoics, Descartes, and Spinoza) on which self-knowledge is attainable and important, although at different times, the focus has been on different forms of self-knowledge (e.g., knowledge of one's nature, character, powers, or psychological attributes.) This hopeful view is the target of skeptics like Pierre Nicole and La Rochefoucauld. See also Garrett (2013).

²¹ I have presented versions of this material at the University of Oxford, University of Pittsburgh, University of South Florida, University of Warwick, and Washington and Jefferson College. Thanks to audiences on those occasions. Thanks also to Lucy Campbell, John Gibbons, Anil Gomes, David Lee, Edouard Machery, Stephen Makin, John McDowell, Joseph Milburn, Matthew Parrott, Karl Schafer, Kieran Setiya, and Ben Sorgiovanni for discussion and written comments, as well as an anonymous referee for this journal.

REFERENCES

- Anscombe, G. E. M. (2000). *Intention*. Cambridge, MA: Harvard University Press.
- Bar-On, D. (2004). *Speaking My Mind*. Oxford: Oxford University Press.
- Boyle, M. (2011). 'Transparent self-knowledge,' *Proceedings of the Aristotelian Society, Supplementary Volume* 85(1), pp. 223–241.
- Burge, T. (1996). 'Our entitlement to self-knowledge,' *Proceedings of the Aristotelian Society* 96(1), pp. 91–116.
- Coliva, A. (2012). 'One variety of self-knowledge: constitutivism as constructivism,' in A. Coliva (ed.) *The Self and Self-knowledge*. Oxford: Oxford University Press, pp. 212–242.
- Davidson, D. (1984). 'First-person authority,' *Dialectica* 38(2/3), pp. 101–111.
- Doyle, C. (forthcoming). 'Internalism and pessimism,' *Journal of Moral Philosophy*. <https://doi.org/10.1163/17455243-20180002>
- Evans, G. (1982). *The Varieties of Reference*, J. McDowell (ed.) Oxford: Clarendon Press.
- Fernandez, J. (2013). *Transparent Minds*. Oxford: Oxford University Press.
- Freud, S. (1966). *Introductory Lectures on Psycho-analysis*. New York: W.W. Norton & Co.
- Gallois, A. (1999). *The World Without, the Mind Within*. Cambridge: Cambridge University Press.
- Garrett, A. (2013). 'Seventeenth-century moral philosophy: self-help, self-knowledge, and the Devil's Mountain,' in R. Crisp (ed.) *The Oxford Handbook of the History of Ethics*. Oxford: Oxford University Press, pp. 229–279.
- Garrett, A. (2017). 'Self-knowledge and self-deception in modern moral philosophy,' in U. Renz (ed.) *Self-knowledge: A History*. Oxford: Oxford University Press, pp. 164–181.
- Gaylin, W. (2000). *Talk Is Not Enough: How Psychotherapy Really Works*. New York: Little, Brown and Co.
- Gibbons, J. (2013). *The Norm of Belief*. Oxford: Oxford University Press.
- Harcourt, E. (2017). 'Containment and 'rational health': Moran and psychoanalysis,' *European Journal of Philosophy*. <https://doi.org/10.1111/ejop.12316>.
- Heal, J. (1994). 'Moore's paradox: a Wittgensteinian approach,' *Mind* 103, pp. 5–24.
- Hieronymi, P. (2009). 'Two kinds of agency,' in L. O'Brien and M. Soteriou (eds) *Mental Actions*. Oxford: Oxford University Press.
- Hills, A. (2009). 'Moral testimony and moral epistemology,' *Ethics* 120(1), pp. 94–127.
- Hills, A. (2010). *The Beloved Self*. Oxford: Oxford University Press.
- Hills, A. (2013). 'Moral testimony,' *Philosophy Compass*, 8(6), 552–559.

- Hopkins, R. (2007). 'What is wrong with moral testimony?' *Philosophy and Phenomenological Research* 74, pp. 611–634.
- Howell, R. (2014). 'Google morals, virtue, and the asymmetry of deference,' *Nous* 48(3), pp. 389–415.
- Lear, J. (2004). 'Avowal and unfreedom,' *Philosophy and Phenomenological Research* 69(2), pp. 448–454.
- Leite, A. (2018). 'Changing one's mind: self-conscious belief and rational endorsement,' *Philosophy and Phenomenological Research* 97(1), pp. 150–171.
- Marcus, E. (2016). 'To believe is to know that you believe,' *dialectica* 70(3), pp. 375–405.
- Martin, M. (1998). 'An eye directed outward,' in C. Wright, B. Smith and C. Macdonald (eds) *Knowing Our Own Minds*. Oxford: Oxford University Press, pp. 99–121.
- McGrath, S. (2009). 'The puzzle of pure moral deference,' *Philosophical Perspectives* 23(1), pp. 321–344.
- Mogensen, A. (2017). 'Moral testimony pessimism and the uncertain value of authenticity,' *Philosophy and Phenomenological Research* 95(2), pp. 261–284.
- Moran, R. (2001). *Authority and Estrangement*. Princeton, NJ: Princeton University Press.
- Nickel, P. (2001). 'Moral testimony and its authority,' *Ethical Theory and Moral Practice* 4(3), pp. 253–266.
- Parrott, M. (2017). 'Self-blindness and self-knowledge,' *Philosopher's Imprint* 16, 1–22.
- Peacocke, C. (1998). 'Conscious attitudes, attention, and self-knowledge,' in C. Wright, B. Smith and C. MacDonald (eds) *Knowing Our Own Minds*. Oxford: Oxford University Press, pp. 63–97.
- Raz, J. (2010). 'The guise of the good,' in S. Tenenbaum (ed.) *Desire, Practical Reason, and the Good*. Oxford: Oxford University Press, pp. 111–136.
- Richfield, J. (1954). 'An analysis of the concept of insight,' *Psychoanalytic Quarterly* 23, pp. 390–408.
- Rödl, S. (2007). *Self-Consciousness*. Cambridge, MA: Harvard University Press.
- Roessler, J. (2013). 'The silence of self-knowledge,' *Philosophical Explorations* 16, pp. 1–17.
- Shoemaker, S. (1994). *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.
- Shoemaker, S. (2012) 'Self-intimation and second-order belief', in D. Smithies and D. Stoljar (eds), *Introspection and Consciousness*. Oxford: Oxford University Press. 239–258.
- Travis, C. (2012). 'While under the influence, in his *Perception*. Oxford: Oxford University Press, pp. 391–412.
- Wallerstein, R. (1995). *The Talking Cures*. New Haven: Yale University Press.
- Whiting, D. (2015). 'The glass is half empty: a new argument for pessimism about aesthetic testimony,' *British Journal of Aesthetics* 55(1), pp. 91–107.
- Wollheim, R. (2003). 'On the Freudian unconscious,' *Proceedings and Addresses of the American Philosophical Association* 77(2), pp. 23–35.