

The Ubiquity of State-Given Reasons^{*}

forthcoming in *Ethics*

Abstract

This paper offers a general argument against the orthodox ‘object-given’/‘state-given’ theory, according to which the distinction between ‘right’ and ‘wrong’ kinds of reason is just the distinction between ‘object-given’ and ‘state-given’ reasons, on the grounds that it fails to account for ‘right’ kind reasons against intention and belief.

Philosophers have come to distinguish between ‘right’ and ‘wrong’ kinds of reasons for belief, intention, and other attitudes. Several theories about the nature of this distinction have been offered, by far the most prevalent of which is the idea that it is, at bottom, the very same as the distinction between what are known as ‘object-given’ and ‘state-given’ reasons. This paper argues that the object-given/state-given theory overgeneralizes on a small set of data points, and in particular that any adequate account of the distinction between the ‘right’ and ‘wrong’ kinds of reason must allow state-given reasons to be of the right kind.

The paper has three main goals. In part A I set up the problem by introducing the right-kind/wrong-kind distinction, the object-given/state-given distinction, and what I will refer to as the object-given/state-given *theory*, according to which the former distinction simply amounts to the latter. Constructing the right way of thinking about the problem constitutes the first main goal of the paper. Parts B and C pursue my second goal by presenting the main argument of the paper: in part B I argue against the object-given/state-given theory by showing that all of the earmarks of the ‘right’ and ‘wrong’ kinds of reason apply to reasons *not* to intend and not to believe, but that these cases can’t be captured by the object-given/state-given theory. And in part C I consider and reject several natural and common responses to the argument in part B. Finally, in part D I tackle my third goal by using the arguments of part B to motivate and explore a more general hypothesis about the right-kind/wrong-kind distinction, and explore some of the consequences of rejecting the object-given/state-given theory.

^{*} Special thanks to Jake Ross, Wlodek Rabinowicz, Ralph Wedgwood, John Broome, Toni Ronnow-Rasmussen, Christian Piller, David Owens, Don Hubin, Simon Kirchin, Jonathan Dancy, Pamela Hieronymi, David Enoch, Stew Cohen, Daniel Star, and to audiences at Oxford, Kent, York, Sheffield, Lund, the University of Texas at Austin, and Boston University, as well as to the helpful referees and editors at *Ethics*. This paper was written with the support of a grant from the Advancing Scholarship in the Humanities and Social Sciences initiative at USC.

A Setting up the Problem

1 *The Right/Wrong Distinction*

The distinction between the ‘right’ and ‘wrong’ kinds of reason is most familiar from the case of reasons for belief, as the distinction between *epistemic* and *pragmatic* reasons.¹ Famously, for example, though both Anselm and Pascal attempted to articulate reasons to believe in God, they took very different strategies. The reasons Anselm sought to articulate at least purported to be *evidence* of God, but not so for Pascal; the reasons Pascal sought to articulate were only *pragmatic* – a kind of penalty or cost to unbelief, or expected reward for belief. This difference – between what are typically called *epistemic* vs. *pragmatic* reasons for belief – is a paradigm of the more general distinction that has come to be known as the distinction between the ‘right’ and ‘wrong’ kinds of reason.

As we’ll see in a moment, the labels ‘right’ and ‘wrong’ are useful because the distinction between epistemic and pragmatic reasons for belief is matched by very similar distinctions among reasons for intention, admiration, amusement, fear, and other attitudes. To see what makes all of these distinctions similar and worth highlighting together under the category of a single ‘right-kind’/‘wrong-kind’ distinction, it helps to start by highlighting a few of the most striking features which do, pre-theoretically, characterize the difference between epistemic and pragmatic reasons for belief. I will call these the *earmarks* of the right-kind/wrong-kind distinction, and they’ll also come in handy later on.

For example, one of the first things that philosophers – including Pascal – noticed about pragmatic reasons for belief, is that they exhibit what I’ll call an *asymmetry of motivation* from epistemic reasons for belief. What I mean by that, is that it is at least much easier to believe in God on the basis of the reasons that Anselm provides – if you find them convincing – than to believe in God on the basis of the reasons that Pascal provides. If you are convinced by Pascal’s argument, what you can do, is take *indirect* strategies: you can spend more time taking the Eucharist, for example, and less time discussing theology with your atheist friends. But it is at least substantially more difficult to believe in God directly on the basis of the reasons that Pascal provides. Some philosophers say that it is *impossible* to believe in God directly on the basis of the reasons that Pascal provides, but that is more controversy than we need to engage in, here.² For my purposes all that we will need is that there is some substantial asymmetry in how straightforward or easy it is to believe on the basis of pragmatic versus epistemic reasons.

A second important difference between the reasons Pascal attempted to articulate and those that Anselm attempted to articulate is that being aware of Pascal’s reasons – even if you grant the success of his argument – does not seem to make belief in God more rational *qua* belief. Whereas being aware of Anselm’s

reasons – at least, if you grant the success of his argument – does seem to make belief in God more rational *qua* belief. So there seems to be a distinctive dimension of rational assessment of beliefs – sometimes called *epistemic* rationality – that is affected by the epistemic reasons of which the subject is aware, but not affected by the pragmatic reasons of which the subject is aware. The same observation goes, whether we are talking about the rationality *of* believing or rationality *in* believing – the distinction that epistemologists sometimes call the distinction between *propositional* and *doxastic* justification. Pascal’s reasons no more affect whether someone who believes in God can be said to do so rationally, than they affect whether it would be rational *for* someone to believe in God.

Focusing on *epistemic* rationality – this distinct dimension of the rational assessment of beliefs – allows us to elide the question of whether there is also some sense in which Pascal showed belief in God to be rational – perhaps a more global or practical sense less central to epistemology. In what follows, though I will drop this qualification and talk simply about the rationality of belief, everything I say should be understood as neutral about whether there is a more inclusive way of assessing the rationality of belief. What is important is that there is some central dimension of rational assessment that is not affected by Pascalian considerations.

A third important difference between epistemic and pragmatic reasons for belief is that epistemic, but not pragmatic, reasons appear to bear on the *correctness* of belief. A belief is correct just in case it is true, and epistemic reasons for belief bear on whether that belief is true, but pragmatic reasons are irrelevant to its truth. So in addition to being subject to asymmetry of motivation and to differentially bearing on the rationality of belief, epistemic reasons differ from pragmatic reasons in bearing on the standards of correctness for belief. Finally, as we’ll see once we compare to some other examples, a fourth feature of the distinction between epistemic and pragmatic reasons for belief is that pragmatic reasons for belief have a recognizable ‘flavor’ that makes them feel intuitively like reasons for other attitudes that exhibit some of the other characteristics of pragmatic reasons for belief.³

So why am I putting so much energy into making explicit these differences between epistemic and pragmatic reasons? It is because it turns out that a strikingly similar distinction arises between reasons in other domains – between reasons which exhibit an asymmetry of motivation, which differ in whether they bear on a distinctive dimension of rationality, and which differ in whether they bear on correctness. Perhaps the most prominent example is the distinction between ‘right’ and ‘wrong’ kinds of reason to intend, as exemplified in Kavka’s famous toxin puzzle.⁴ In the toxin puzzle, you are offered a reward for forming an intention to do something that will have a bad result – a reward that attaches to the intention, not to the action. About the toxin puzzle, philosophers have observed (I) an asymmetry of motivation – it

is harder to form an intention on the basis of rewards you get for having the intention, than to form an intention on the basis of rewards for doing the action; (2) asymmetric effects on rationality – toxin-puzzle-like reasons seem not to affect the rationality of having the intention – neither the rationality of forming the intention, nor anyone’s rationality *in* having the intention; and (3) the incentive to form the intention seems to have the same ‘flavor’ as pragmatic reasons for belief – it feels like something similar is going on. (Though it is more controversial whether toxin-puzzle-like reasons to intend bear on the correctness of intention, that is simply because it is more controversial what it takes for an intention to be correct than for a belief to be correct.)

So by and large, toxin-puzzle-like reasons bear the same earmarks as pragmatic reasons for belief such as those which Pascal attempted to provide. Consequently, many philosophers have come to think about the relationship that toxin-puzzle-like reasons bear to more ordinary reasons for intention as the *same in kind* as the relationship that pragmatic reasons bear to epistemic reasons for belief.⁵ This relationship has come to be known as the distinction between ‘wrong’ and ‘right’ kinds of reason, and the same distinction has been observed to arise in many other cases, as well – between ‘right’ and ‘wrong’ kinds of reasons to fear, admire, blame, be amused, and more, for which it is marked out by the same earmarks as for belief and intention.⁶

The question of what distinguishes the ‘right’ kind of reason from the ‘wrong’ kind turns out to be interesting for a variety of reasons.⁷ At the highest level of generality, independent of the particular things that make this question interesting, the ‘wrong kind of reasons problem’ is the problem of saying just what this distinction amounts to – what *makes* some things ‘right-kind’ reasons and other things not. Note that this question arises equally for views according to which so-called ‘wrong-kind reasons’ aren’t even reasons at all, strictly speaking, or are only reasons to do something else.⁸ Even such a view must be able to tell us what sort of things turn out to genuinely be reasons for belief (or for intention), as compared to the things which aren’t even reasons for belief or intention at all.

2 *The Object-Given/State-Given Distinction*

The most dominant theory about the nature of the distinction between the ‘right’ and ‘wrong’ kinds of reasons attributes that distinction to the difference between what are called often ‘object-given’ and ‘state-given’ reasons. Though both terminology and details differ between different proponents, the basic idea behind each attempt to formulate a distinction along these lines stems from trying to generalize on an observation about which reasons seem to be the right kind of reasons for belief, and which reasons seem to be the right kind of reasons for intention.

The observation goes like this: intuitively, the ‘right kind’ reasons to believe that p are just evidence that p . So since the proposition that p is the object of the belief that p , right-kind reasons for belief bear on the objects of the beliefs they are reasons for. Similarly, intuitively, what goes wrong in the toxin puzzle is that the incentive to intend to drink the toxin does not come with any corresponding reason to actually drink the toxin. So since *drinking the toxin* is the object of the intention to drink the toxin, right-kind reasons for intention must be ones that bear on the objects of the intentions they are reasons for. So the idea that right-kind reasons differ from wrong-kind reasons by bearing on the objects of the attitudes they are reasons to have, is an attempt to generalize on the idea that the right-kind reasons for belief include all and only the evidence, and the idea that right-kind reasons for intending to do A include all and only the reasons for and against doing A. As I understand the various versions of the object-given/state-given theory, they all aspire to say something general that captures these two insights as special cases by requiring that right-kind reasons for some attitude must ‘bear on the object’ of that attitude, while wrong-kind reasons just bear on the benefits of being in the state of having that attitude (and hence count as ‘state-given’).⁹

It turns out that saying precisely what it means to require that a right-kind reason must ‘bear on the object’ of an attitude is much harder than many people have let on. Here is one important reason why: according to a natural view about the objects of intentions, they are propositions. When you intend to do A, the object of your intention is the proposition that you do A. Not everyone accepts this view, but the evidence for it is the same as the evidence that desires take propositions for their objects. Just as you can want to go to the beach, you can want Jack to go to the beach. So rather than assuming that ‘want’ is polysemous or ambiguous between these two constructions, linguists assume that ‘want’ has the same meaning and syntactic structure, but that in ‘you want to go to the beach’, there is a silent pronoun between ‘want’ and ‘to’. On this view, for which there is additional evidence from syntax, ‘you want to go to the beach’ means approximately, ‘you want yourself to go to the beach’ – that’s why we take desires (wants) to have propositional objects, even when they are desires about what to do.¹⁰ And the same evidence goes for intention: just as you can intend to go to the beach, you can intend Jack to go to the beach, or intend for Jack to go to the beach. This can be easily explained by the hypothesis that ‘intend’ takes propositional arguments, and when you intend to go to the beach, what you intend is for *yourself* to go to the beach, rather than intending for someone else to go to the beach.

So much for our digression on the objects of intention. The problem is that if intention and belief take the same objects, there is nothing in the loose formulation of the object-given/state-given theory that we’ve given so far that explains, for example, why evidence that you will go to the beach is not a right-kind

reason for you to intend to go to the beach: after all, it bears on the proposition that you will go to the beach, and that proposition is the object of the intention to go to the beach. Similarly, there is nothing in the loose formulation of the object-given/state-given theory that we've given so far that explains why reasons for you to go to the beach are not *ipso facto* right-kind reasons for you to believe that you will go to the beach. But if the former of these is true, then right-kind reasons for intending to do A are not limited to reasons to do A after all, and if the latter is true, then right-kind reasons for believing that *p* are not limited to evidence that *p* after all.

Moreover, since the right-kind/wrong-kind distinction is supposed to apply to reasons for a wide variety of attitudes, this problem is not limited to the case of belief and intention, and does not require the assumption that belief and intention may have the same objects. If *any* pair of attitudes have the same type of objects – wanting, hoping, fearing, or any of the rest – then we will be able to construct a similar problem for such a pair. This problem makes it actually quite difficult to give any content to the slogan that right-kind reasons must 'bear on the object' of the attitude, because the problem shows that it is not enough to 'bear on the object' of the attitude; a right-kind reason must also bear on that object *in the right way*.¹¹ But then we still need an account of what the right way is.

There are also other problems about formulating the object-given/state-given distinction correctly, some of which are familiar from the literature.¹² Fortunately, however, we need not try to solve those problems here, because given the motivations that I described earlier, I take it that any attempt to successfully formulate a general version of the object-given/state-given theory is an attempt to find a generic formulation which predicts **believe object** and **intend object** as special cases:

believe object R is a right-kind reason bearing on believing *p* just in case R is evidence for or against *p*.

intend object R is a right-kind reason bearing on intending to do A just in case R is a reason bearing on whether to do A.

So rather than worrying about whether there is any neat way of formulating such a theory, I'll be arguing against **believe object** and **intend object** directly. Hence, if my arguments work, they will work against any way of making this theory precise, whether it uses terms like 'object-given' or not.¹³

3 *An Earlier Argument*

In previous work, I argued against the object-given/state-given theory on the grounds that it is insufficiently general.¹⁴ Appealing to what I've here called the third and fourth earmarks of the right-

kind/wrong-kind distinction – the fact that only right-kind reasons bear on standards of correctness, and the intuitive ‘flavor’ of wrong-kind reasons – I argued that the right-kind/wrong-kind distinction arises for things other than mental states or attitudes – and hence for cases in which we can’t even distinguish between a ‘state’ and its ‘object’. For example, I claimed that because there are correct and incorrect ways to tie a knot, there are right and wrong kinds of reasons to manipulate rope in certain ways – with respect to the activity of tying a knot. Being offered ten dollars to manipulate a rope in certain ways does not make those the correct way to tie a knot, and has the intuitive flavor of ‘wrong-kind’ reasons more generally. But manipulating rope is not a mental state, and does not have an object. So we can’t distinguish between reasons which bear on its object and those which do not. Hence, I claimed, the object-given/state-given distinction is too narrow to account for whatever is going on in the distinction between ‘right-kind’ and ‘wrong-kind’ reasons to tie a knot in certain ways.

I still believe that the main idea of this earlier argument is on the right track, for the object-given/state-given theory is indeed too narrow. But the knot-tying argument has a major weakness. For the kinds of cases that it appeals to do not exhibit all of the earmarks of the right-kind/wrong-kind distinction. In particular, there is no difficulty in being motivated to manipulate rope in ways that are incorrect ways to tie a knot, for the reason that you have been offered ten dollars to do so, and there may be no irrationality in doing so – if having a secure knot is worth less than ten dollars, at any rate. The argument also leaves open whether the object-given/state-given theory is at least extensionally correct about the right and wrong kinds of reasons for and against attitudes like belief and intention, and hence whether **believe object** and **intend object** are true. Consequently, it is easy to see why a proponent of the object-given/state-given theory might find this argument quite unconvincing.

In contrast, my argument in this paper attacks the object-given/state-given theory on its home territory. Like my earlier argument, the argument of this paper will work by trying to show that the right-kind/wrong-kind distinction applies to reasons that are not reasons for attitudes with objects at all – which is why the object-given/state-given distinction can’t apply to them. But what I will be arguing is that the object-given/state-given theory is not even extensionally correct *even for its home territory*: belief and intention. And just as significantly, I will be arguing this by appeal to the ‘core’ earmarks of the right-kind/wrong-kind distinction: asymmetry of motivation and bearing on rationality – again in contrast to my earlier argument.

B The Argument

1 *Reasons Against Intending*

This section contains the core ideas of the entire paper. What I'll show in this section is that there are considerations which bear all of the marks of right-kind reasons *against* intending to do A, which do not themselves bear on whether to do A. The argument is simple: if these considerations bear all of the marks of right-kind reasons, they *are* right-kind reasons – after all, the 'right-kind'/'wrong-kind' distinction was just a catch-all label designed to cover an important class of differences that arise in a variety of domains. But since **intend object** classifies them as being of the wrong kind, any theory which is committed to **intend object** – and in particular, any version of the object-given/state-given theory – is false. Section B2 extends this same pattern of argument to reasons against belief and section B3 returns to defend my characterization of the object-given/state-given theory and extend the argument. Part C will take up several important lines of possible resistance.

The cases in which I will be interested are cases in which it is either rationally permissible or rationally required neither to intend to do A, nor to intend to not do A, but to form neither intention, at least for the interim. The structural feature of intention which makes this class of counterexamples possible is that while there is no third option other than doing A and not doing A, there *is* a third possibility besides intending to do A and intending to not do A – it is having neither intention.¹⁵ The central flaw in the object-given/state-given theory is that it ties right-kind reasons to intend to do A too closely to reasons to do A, and right-kind reasons to intend to not do A too closely to reasons to not do A, leaving no room left over for right-kind reasons to have neither intention. But as I'll be arguing, it certainly *looks* like there can be right-kind reasons to have neither intention.

I'll focus on two different kinds of cases in which it is rational to not form an intention about whether to do A, although I believe that there are more. The first case is one in which the agent expects further information to come to light, before a decision needs to be made. For concreteness, suppose that I am trying to decide whether to drive to Los Angeles tomorrow morning. One of the benefits of doing so is that I get to enjoy the Pacific coastline on the drive down. But one of the costs is that I am likely to get stuck in traffic. There are other potential benefits and costs, but I don't know what they all are. One potential benefit, however, is that I might get to see my brother. If he is in LA, then that would definitely make it worthwhile to make the trip, but if he is not in LA, then the scenery will not be worth the traffic.

Currently I think he is likely enough not to be there, that it's not worth the traffic. So if I have to make a decision right now, the best choice, given what I know, is to stay home. But fortunately, I'm

expecting my brother to call later this afternoon to let me know whether he will be in LA or not. I'm expecting him to tell me that he won't, of course, but that's just my estimate – I also expect to be better informed after he calls, and it is quite possible that he will tell me that he *will* be there. And if he is there, that would make the trip worthwhile. In a situation like this one, it makes perfect sense for me to wait to decide whether to drive to LA tomorrow.

In fact, given my belief that later this afternoon – in plenty of time before I need to make preparations to leave in the morning – I will have more information that has the potential to make a difference in what I would decide, and given that I don't need to make any decisions in the meantime that depend on my plans for tomorrow – such as coordinating with my wife who gets the car – it seems to me that waiting to decide is the *only* rational course. It is not only rational *for* me to form neither intention now; if I do take the fact that my brother will call me later this afternoon into account and wait to decide for that reason, then I am rational *in* forming neither intention now. And there is no intuitive difficulty in forming neither intention on the grounds that more information will soon come to light – on the contrary, it is easy to wait to decide for exactly this sort of reason. So my knowledge that more relevant information will come to light has the key earmarks of a right-kind reason to lack intention.

In the case as I've described it, I've only considered three options – intending to drive into LA, intending not to drive into LA, and having neither intention. It is also possible, of course, that I could form a conditional intention to drive in if I hear from my brother that he will be there, and to stay home otherwise.¹⁶ This fourth option may be at least as advisable as not having any intention at all, but its existence does not interfere with my main observation about the case, which is that the fact that further information is forthcoming counts *against* my having an unconditional intention now to stay home. Moreover, though conditional intentions are important in some simple cases, in more complex cases in which more information is forthcoming, it may be more advisable to wait on that information, than to form a plan in advance about how to act, conditional on it.

In a second kind of case, the agent is making a complicated decision which requires some thought, but is confronted by more pressing current challenges. For example, suppose that I am deciding which graduate school to attend. I receive offers of admission from Old-School Eastern and Western Powerhouse in mid-February, and have until April 15th to make my decision. The decision is one of the most important ones for my career and lifestyle, and will affect me for years to come. But in the meantime, I'm talking to my mother-in-law on the phone, my wife is reminding me (for the second time) to take out the trash, and my smoke alarm is going off. Since I have plenty of time to make the decision, the rational thing for me to do is to take care of my smoke alarm (and the trash) first, and trust myself to make the

decision about whether to go to Old-School Eastern or Western Powerhouse later, when I have more time. Again, it is perfectly possible to put off this decision on these grounds – we don't observe any asymmetry of motivation. And these circumstances *do* affect whether it is rational to form the intention to go to Old-School Eastern or Western Powerhouse, at this time – the rational thing to do is to put off the decision and take care of the smoke alarm. So the smoke alarm's going off right now has the earmarks of a right-kind reason against forming either intention about where to go to graduate school.

Both of these cases involve a consideration which exhibits the marks of a right-kind reason against forming an intention – a right-kind reason not to intend. But neither of them involves an object-given reason, as the object-given/state-given theory would require. The fact that my brother will call me later is not a reason which bears either way on whether I should drive to LA tomorrow. So it doesn't bear on the object of the intention to drive to LA tomorrow or on the object of the intention to not drive to LA tomorrow. But as we've seen, it bears on these intentions nonetheless, because being aware of it can make it rational – even uniquely rational – to have neither one. Similarly, the fact that my smoke alarm is going off is not a reason which bears either way on whether to go to Old-School Eastern or Western Powerhouse for graduate school. So it doesn't bear on the object of the intention to go to Old-School Eastern or on the object of the intention to go to Western Powerhouse. But as we've seen, it bears on these intentions nonetheless, because being aware of it can make it rational – even uniquely rational – to form neither one.

Moreover, and importantly, the two cases of earmark-bearing right-kind reasons not to intend that we have just been considering contrast sharply with intuitively clear cases of wrong-kind reasons not to intend. For example, suppose that I am offered one thousand dollars to not make up my mind about whether to wear clothes or go naked to the talk I am giving tomorrow in the Oxford moral philosophy colloquium series until thirty minutes before I am scheduled to begin. Although I can imagine taking indirect strategies toward achieving this – by spending the time on the phone with my nudist friends, for example, and distracting myself from thinking about the look I expect to appear on John Broome's face if I show up in the buff, it is not at all an easy thing to do, simply for the prospect of money – and not simply because the financial award isn't large enough. Similarly, though lacking an intention about whether to wear clothes to the talk would be an advantageous state of mind for me to be in, I still think that when we restrict ourselves to the distinctive rational evaluation of intention on analogy with the evaluation of beliefs as *epistemically* rational, the only rational state, *qua intention*, for me to be in, is to intend to wear clothes. So the problem isn't simply that there *aren't* wrong-kind reasons to lack intention – on the contrary, what this example along with the previous examples shows, is that the right-kind/wrong-kind distinction arises just as much for *lacking* intention, as for *having* intention.

The argument in this section – the central argument of this paper – has been a simple one. The argument is that there are considerations which stand to lacking intention exactly as we would expect right-kind reasons to do. They are reasons for which it is perfectly possible and natural to lack intention, awareness of them tends to make it rational *for* you to lack intention, and when you lack intention on the basis of them, you are rational *in* lacking intention. Moreover, they do not have the intuitive flavor exhibited by other wrong-kind reasons. The only earmark of right-kind reasons which I haven't focused on, is the way in which they bear on standards of correctness, but as before, since the standards of correctness on intention are controversial, this simply complicates matters. But these reasons are *not* object-given, and they *are* intuitively state-given, because they bear on the benefits or costs of forming the intention – in the first case, forming the intention closes a deliberative question that it may turn out to be best to leave open, and in the second case, forming the intention distracts from something else that is more pressing. So these considerations appear not only to be right-kind reasons against intention, but to certainly not be object-given, and intuitively to be state-given.

I'll close this section by briefly revisiting the toxin puzzle, which plays a central role in motivating the object-given/state-given theory, by motivating **intend object**. We are now in a position to observe that **intend object** is really a significant overgeneralization on what we observe in the toxin puzzle case, because in the toxin puzzle case, we only observe a reason *to* intend which is of the wrong kind – but our examples here have focused on reasons *not* to intend. In other words, it is crucially important to distinguish between **intend object** and **intend okay**:

intend object R is a right-kind reason bearing on intending to do A just in case R is a reason bearing on whether to do A.

intend okay R is a right-kind reason to intend to do A just in case R is a reason to do A.

If we substitute 'not-A' for 'A' in **intend okay**, we get

intend okay* R is a right-kind reason to intend to do not-A just in case R is a reason to do not-A.

The problem that we have been observing, is that so long as the reasons bearing on whether to do A are exhausted by the reasons to do A and the reasons to do not-A, it is a consequence of **intend object** that R is a right-kind reason bearing on the intention to do A just in case R is either a reason to intend to do A or a reason to intend to do not-A. Hence, **intend object** goes above and beyond **intend okay** by ruling out

right-kind reasons to lack intention. But this is *not* a consequence of **intend okay** itself. So everything that we have said in this section is compatible with **intend okay**.¹⁷

Consequently, it is highly natural to see **intend object** as a hasty overgeneralization on the much more plausible **intend okay**, which is much more directly motivated by the actual observations that it is possible to make about the toxin puzzle. This allows us to see how the object-given/state-given theory could have gone so wrong. It is a simple oversight which makes sense in the light of a selective palette of examples.

2 *Reasons Against Belief*

The same kinds of considerations as those in the last section show that as with intention, there are right-kind reasons against belief that are not object-given, and intuitively are state-given. Again, this result turns on the fact that the reasons which bear on whether p are exhausted by evidence for p , which are right-kind reasons to believe that p , and evidence for not- p , which are right-kind reasons to believe not- p . But believing p and believing not- p are not the only options, when it comes to belief; a third option is always to lack any belief with respect to p .¹⁸ And in fact often the most rational option is to be unopinionated with respect to p – particularly whenever the evidence for p and for $\sim p$ is relatively equally balanced. But since whatever considerations make it rational to lack belief are neither evidence for p nor evidence for $\sim p$, the object-given/state-given theory would seem to predict that they cannot be of the right kind.¹⁹

Again, I believe that there are several kinds of examples which illustrate this general point, but again, one of the clearest examples comes from cases in which the agent expects further evidence to be forthcoming. For example, suppose that several times in the past my doctor has found spots on my skin that are suspicious for cancer, but that each time biopsy has revealed them to be benign. Once again, my doctor finds a suspicious spot on my skin, and takes a tissue sample to send to the lab. Based on inductive evidence, the tissue is likely to be benign, but it makes sense to wait on the test results before concluding that it is really benign. The fact that the lab is soon going to report its results isn't evidence that my skin spot is benign, and it isn't evidence that it is not benign, but it is a reason to not yet make up my mind as to whether it is benign. It is easy to not yet make up my mind on the basis of this reason, and it makes it rational to not yet make up my mind. It has the earmarks of a right-kind reason, but it isn't object-given.

Cases like this one contrast sharply with cases of wrong-kind reasons to not yet make up my mind. For example, suppose that I am offered one thousand dollars to not make up my mind about some question for which all of the relevant evidence has already come in. Though it's easy to not make up my mind on the grounds that further evidence is forthcoming, it's hard to not make up my mind directly on

the grounds that I'll get money for having not made up my mind – though of course, there are always indirect strategies, like paying close attention to some evidence and trying to ignore the counterevidence. So as with intention, there are contrasting cases of right-kind and wrong-kind reasons to lack belief – to have not made up your mind.

The same points about the reason for this failure go for intention as for belief. I'll close this section by briefly revisiting Pascal's Wager, which plays a central role in motivating the object-given/state-given theory, by motivating **believe object**. We are now in a position to observe that **believe object** is really a massive overgeneralization on what we observe in Pascal's Wager, because in Pascal's Wager, we only observe a reason *to* believe which is of the wrong kind – but our examples here have focused on reasons *not* to believe. In other words, it is crucially important to distinguish between **believe object** and **believe okay**:

believe object R is a right-kind reason bearing on believing p just in case R is evidence for or against p .

believe okay R is a right-kind reason to believe p just in case R is evidence for p .

If we substitute ' $\sim p$ ' for ' p ' in **believe okay**, we get

believe okay* R is a right-kind reason to believe $\sim p$ just in case R is evidence for $\sim p$.

The problem that we have been observing, is that **believe object** goes above and beyond **believe okay** by ruling out right-kind reasons to lack belief. But this is *not* a consequence of **believe okay** itself. So everything that we have said in this section is compatible with **believe okay**.²⁰

Consequently, it is highly natural to see **believe object** as a hasty overgeneralization on the much more plausible **believe okay**, which is much more directly motivated by the actual observations that it is possible to make about Pascal's Wager. This allows us to see once more how the object-given/state-given theory could have gone so wrong. It is a simple oversight which makes sense in the light of a selective palette of examples.

3 *Intend Okay and Believe Okay*

At this point careful readers will wonder whether I haven't made things too easy for myself by characterizing the object-given/state-given theory as committed to **intend object** and **believe object**, as opposed to characterizing it as simply committed to **intend okay** and **believe okay**. After all, in my

original introduction of the object-given/state-given theory, it was motivated by observations about wrong-kind reasons *to intend* and *to believe*. And all that those cases directly give us, are **intend okay** and **believe okay**. To this I have two responses. In order, they are 1) that a commitment to **intend object** and **believe object** is necessary for the object-given/state-given theory to give us a general account of the nature of the right/wrong distinction, and 2) that in any case my argument from section BI can plausibly be generalized – at least against **intend okay**.

It may indeed be that some proponents of the importance of the object-given/state-given distinction or its close cousins have not meant to adopt any commitments about the conditions under which something is a right-kind reason not to intend or not to believe. Such theorists need only be committed to **intend okay** and **believe okay**, and since I have not yet argued that either of these theses is false, I have not yet argued against such views. However, my arguments do show that such a view is too limited to be able to amount to a theory of the *nature* of the distinction between the right and wrong kind of reasons. And I defined the object-given/state-given *theory* as the theory that the right/wrong distinction just *amounts to* the object-given/state-given distinction.

My arguments show that there is a right/wrong distinction on reasons against intention and belief that is precisely parallel to the right/wrong distinction on reasons for belief. There is at least as much ground to think that there is a single right/wrong distinction on the against side and on the favoring side, as there is to think that there is a single right/wrong distinction for both belief and intention. But since the lack of intention or belief is not an attitude with an object, the object-given/state-given distinction cannot apply to it. So it is too limited in structure to point us in the direction of a general account of what is going on. This is what I think is important.

However, once my arguments from the previous sections are on the table, I believe that even the retreat to **intend okay** and **believe okay** is not safe. For in addition to there being right-kind reasons to lack intention that are not object-given, because they are reasons to wait to make up your mind, there are also right-kind reasons to *have* an intention that are not object-given, because they are reasons to make up your mind, without being reasons to make it up in any particular way. Suppose, for example, that in the case in which I am thinking about whether to drive into LA tomorrow, I find out that I must coordinate with my wife who will have the car tomorrow before I hear from my brother. That gives me a reason to make up my mind now about whether to go. And if coordinating with my wife is important enough and I think it is unlikely enough that my brother's phone call will inform me that he will be in LA, then this kind of consideration can make it uniquely rational for me to intend, now, to stay home tomorrow – even though I know that further information is forthcoming.

The fact that I need to coordinate now with my wife who will have the car tomorrow is neither a reason for me to drive into LA, nor a reason for me not to drive into LA. She will do just fine either with the car or without it; she simply needs to plan ahead. So it is not an object-given reason for me not to drive into LA. And it is, intuitively, a state-given reason for me to intend not to drive into LA. Having this intention now has the benefit of allowing my wife to plan her day accordingly.

As before, this reason also *contrasts* with clear cases of wrong-kind reasons for me to make up my mind whether to drive into LA tomorrow. For example, suppose that I am offered one thousand dollars to make up my mind now whether to drive into LA tomorrow, rather than waiting to find out about whether my brother will be there. Let's just specify, for concreteness, that I'm 80% confident that my brother will be in LA tomorrow and hence that he will tell me so this afternoon, but that a merely 80% chance of seeing my brother isn't worth the 100% chance of traffic. Since an 80% chance of my brother's being in LA doesn't make going worth it, if I have to decide now, surely the only rational choice would be to decide not to go. But I'm 80% confident that he will call this afternoon and tell me that he will be in LA. And I know in advance that once I know that he will be in LA, the only rational choice for me would be to decide to go. So I'm 80% confident in advance of forming the intention not to go, that I will be rationally required to change my mind in the future. This is exactly the same sort of situation that makes it difficult to form the intention to drink the toxin, in the toxin puzzle – when you know, in advance, that when the time comes, the sole rational choice will be to not drink it. So it looks like the wrong kind of reason to make up my mind.

In contrast, in the case in which I coordinate with my wife over who gets the car, even though I am 80% confident in advance that my brother will call and tell me that he will be in LA, this does not make me 80% confident that I will be rationally required to change my mind, because in this case, having formed an intention not to drive into LA, I expect to have acted on that intention – by allowing my wife to make plans that depend on her having access to the car – in ways that will then make it irrational to decide to go into LA, even though my brother will be there. So this kind of case does not share the problematic feature of the toxin puzzle case that I can be arbitrarily confident that I will later be rationally required to change my mind.

I've just been arguing that the same kinds of considerations which show that there are right-kind reasons *not* to intend that are not object-given can plausibly be extended to show that there are right-kind reasons *to* intend that are not object-given. Both sorts of reasons, intuitively, concern whether it makes sense to make up your mind. This gives us an argument against even **intend okay**. I suspect, however, that the argument cannot be extended to an argument against **believe okay**. This is because whereas action

requires intention, it does not require belief. It is always possible, at least in principle, to act in ignorance. Consequently, it is either always rationally permissible to remain agnostic, or, insofar as you are rationally required to form a belief, that is solely because of the force of the evidence.

If my foregoing characterization is correct, however, of the commitments the object-given/state-given theory must adopt in order to amount to an account of the right/wrong distinction, then neither extension of the argument is necessary. Either way, no version of the object-given/state-given theory is going to have enough generality to give us an account of the nature of the right/wrong distinction.

C **Objections and Replies**

1 *First Objection: Object-Given Reasons to Withhold*

When I have given the argument in the preceding sections, a common response has been to suggest, first, that the right-kind/wrong-kind distinction that I have located is not really a distinction among reasons *not* to intend or *not* to believe, but rather a distinction among reasons *for withholding* belief or intention. Second, this response goes, *withholding* intention with respect to A is not just a matter of *lacking* intention; it is itself a positive attitude – and similarly for withholding belief. Finally, this response goes, since withholding is an attitude after all, the object-given/state-given theory can apply to reasons for it after all – the right-kind reasons featured in my argument will simply turn out to be reasons which bear on the object of *withholding*, rather than on the object of *belief*. Which only makes sense, the objection goes, because we are talking about reasons for withholding, after all, and not reasons for belief.

There is much that I admire about this response to my argument. It accepts at face value, as I believe that we should, the motivation given by the earmarks for thinking that my examples are genuine examples of a distinction between right-kind and wrong-kind reasons. And it accepts the basic fact that underlies both my present argument and my earlier argument – namely, that the object-given/state-given theory only applies to attitudes that have objects. And finally, it agrees with my assumption that my examples do not involve reasons bearing on whether to do A, or evidence bearing on *p*. Given these three constraints, it makes the only move possible: it hypothesizes that the reasons in my example are reasons for some attitude *other* than belief or intention. And the idea that withholding is itself an attitude conveniently supplies the justification for this move. It's a very neat story.

The first thing that we would like to be able to do, in order to evaluate this view, is to check and see whether my examples really do involve reasons which bear on the object of withholding. Unfortunately, I believe it will be difficult for a proponent of the object-given/state-given theory to get the

right result about this. That is because it must be able to explain why my examples of right-kind reasons not to intend or believe bear on the object of withholding, but my examples of wrong-kind reasons don't. Suppose, for example, that we say that the object of withholding is the intention withheld. Unfortunately, it is easier to see how the fact that I've been offered \$1000 to lack this intention could bear on the object of withholding, so understood, than to see how the fact that more information is forthcoming could. For the former mentions the object of withholding, and the latter does not.

It's not enough, in short, to locate an attitude – call it 'withholding' – and say that my cases involve object-given and state-given reasons for it; you must also be able to say what this attitude's objects are, and explain why my examples track the object-given/state-given distinction, so understood. Importantly, one of the hardest parts of making good on this idea is going to be that if the proponent of this objection is right that there is a genuine distinction between withholding intention or belief and merely lacking it, then I can be careful (as in fact I was) to describe my wrong-kind cases as cases in which you are offered \$1000 to *lack* the requisite intention or belief, rather than as cases in which you are offered the money to *withhold* with respect to it.

I believe that getting the object-given/state-given theory to predict the right results about my cases is a genuine obstacle to getting this strategy to work. But as I noted in section A2, it's hard to make precise sense of how to apply the object-given/state-given theory at the kind of level of generality that would let us apply it to new attitudes like withholding. Though this lack of precision is in general a problem for this view, it's also a problem for being able to show in any conclusive way that it *can't* generate the right predictions, here. (This is a special case of the general principle that vague theories are difficult to refute.) So we should look somewhere else for a more conclusive reason to set this alternative view aside.

Fortunately, there is a more conclusive problem with this view. And it is that the right-kind reasons in my examples don't just make it rational to *withhold* belief or intention; they also make it rational to *lack* belief or intention, even in cases in which this doesn't involve withholding. To see why, compare two cases. In both cases, I have the option of driving to LA tomorrow, know that if I do I will get to enjoy the Pacific coastline but am likely to get stuck in traffic, and know that the chance that my brother will be in LA is high enough to make the traffic worth the risk that he won't, though if I knew that he were not to be in LA, then the traffic would not be worth it.²¹ In the first case, as before, I know that more information is forthcoming, because I expect my brother to call tonight and tell me whether he will be in town. But in the second case, I know that no more information is forthcoming – perhaps because my brother has already told me he won't decide for certain whether to be in town until tomorrow morning. In

neither case, however, do I deliberate about whether to drive into LA tomorrow. I just don't think about it, and so while I *lack* the intention to drive into LA, I am not *withholding* this intention.

Because in these two cases I never deliberate about whether to drive into LA, neither involves a case of withholding. They both simply involve a lack of intention. However, there is an important difference between the cases. In the second case, it would be rational for me to intend to drive into LA tomorrow, because there is no more information forthcoming that could rationally require a change of mind. But in the first case, it would not be rational for me to form this intention. So whereas in the second case my lack of intention is rationally *optional*, in the first case my lack of intention is rationally *required*.

Since neither of these cases involves withholding, but merely involve the bare lack of intention, they show that the effects of the right-kind reasons in my examples are not restricted to a positive attitude of withholding, but count directly against intention. That is why I believe that the reasons I've located are genuinely reasons *not to intend*, rather than reasons for some distinct, positive, attitude of *withholding*. If withholding is really a positive attitude, therefore, it is merely the way that we keep ourselves from intending, once we've begun deliberating.

2 *Second Objection: Two Stages*

Another common response to my argument from part 2, is to appeal to a two-stage theory about the rationality of belief and intention. According to the two-stage theory, whether it is rational to believe p depends on two things: first, whether it is rational to make up one's mind with respect to p , and second, whether p is better-supported than $\sim p$. Proponents of the two-stage theory claim that the reasons in my examples are reasons at the first stage, rather than reasons at the second stage, and suppose that this distinction confounds the significance of my argument.

I agree with the two-stage theory that it is rational to believe p only if it is both rational to make up one's mind with respect to p and p is better-supported than $\sim p$. But I suspect that it is impossible to cleanly separate this into two 'stages'. This is because among the factors which affect whether it is rational to make up one's mind about p are the evidence in favor of and against p – the very considerations which according to the two-stage theory are supposed to play a role at the second stage. If the evidence is too evenly balanced or merely probabilistic in nature, then that can make it irrational to make up one's mind, and if the evidence is preponderant and conclusively one-sided, that can make it rationally permissible to make up one's mind no matter what reasons there might be not to.

But more importantly, the two-stage theory does not have the right structure to help the object-given/state-given theorist. For even if the two-stage theory is correct, my examples still show that there is a

right-kind/wrong-kind distinction at the first stage, which the object-given/state-given theory cannot capture. For there is an important difference between cases in which further information is expected and cases in which the agent has been offered money to refrain from making up her mind. This is the right-kind/wrong-kind distinction at work. By the lights of the two-stage theory, these aren't really reasons for and against belief (which only apply at the latter stage), but just reasons for and against making up one's mind. No matter; since not making up one's mind is not itself an attitude with an object, the object-given/state-given theory cannot successfully distinguish between right-kind and wrong-kind reasons not to make up one's mind. So the object-given/state-given theory still fails to capture the distinction between these right-kind and wrong-kind reasons.

3 *Third Objection: Sufficiency Does the Work*

In this section I want to consider and address what I take to be the most important last-ditch line of defense of the object-given/state-given theory – a defense which grants much of my intuitive data, but attempts to account for them without admitting that there are right-kind *reasons* not to intend.

The key idea of this defense is that lacking intention is not just another alternative to intending and intending not – it is the *default*. After all, infants aren't born with intentions about what they will do in one circumstance or another; they lack intentions about everything. So while for intending to do A to be rational, you have to have sufficient right-kind reasons to intend to do A, and while for intending to not do A to be rational, you have to have sufficient right-kind reasons to intend to not do A, the idea behind this reply to my argument is that for *lacking* intention to be rational, you only need to *lack* conclusive right-kind reasons to intend to do A, and lack conclusive right-kind reasons to intend to not do A. According to this view, the kinds of factors that I've isolated – the fact that new information is forthcoming, that current circumstances are more pressing, and that you'll do it anyway – *do* tend to make it rational to lack an intention, but they don't do so by being *reasons* to lack an intention; they do it by raising the bar on how good reasons to *have* an intention need to be, in order to be *sufficient*. Call this the *primitive sufficiency view*. The primitive sufficiency view allows for a neat treatment of my data about how these factors affect the rationality of intention (and corresponding points go for belief) without admitting that there are actually *reasons* to withhold.

Unfortunately, I believe that there are several problems with the alternative explanation of my data offered by the primitive sufficiency view. The chief of these is that it crucially leaves the notion of *sufficiency* unexplained. In this account, whether or not a reason or set of reasons is sufficient plays a very important

role, but the account doesn't have the resources to say what sufficiency consists in or what makes reasons sufficient – it merely needs to appeal to sufficiency as a *deus ex machina* to do the right sort of work.

This problem, in turn, is aggravated by a second problem: that even the proponent of the primitive sufficiency view must acknowledge that the sufficiency of reasons is a complicated matter, and is affected by a range of different factors – but not affected by other factors. For example, it takes more for reasons to intend to be sufficient when further information is forthcoming, but it doesn't take more for them to be sufficient when you are offered money to not make up your mind. Something has to distinguish between the 'right kind' of factors to raise the bar on what it takes for reasons to be sufficient, and the 'wrong kind' of factors to raise this bar. The fact that the bar for sufficiency behaves in this complicated way aggravates the first problem by making sufficiency a particularly problematic thing to take as a primitive, and constitutes a problem in its own right, because the contrast between the 'right' and 'wrong' kinds of factors to raise the bar on sufficiency looks an awful lot like the problem of distinguishing between the 'right' and 'wrong' kinds of reasons with which we began.

A third problem with the primitive sufficiency view is that though it explains how factors could make it harder for intention to be rational without themselves being reasons not to intend, it doesn't explain how we could lack intention *for* these reasons. But as I argued in section BI, it *does* seem that I can put off forming an intention about whether to drive to LA on the grounds that I will soon find out whether my brother will be there or not – and in particular, this *contrasts* with my ability to put off forming an intention about whether to wear clothes to my talk on the grounds that I'll be financially rewarded for not making that decision until within half an hour of the talk. If the factors that raise the bar on the sufficiency of reasons to intend are not themselves right-kind reasons not to intend, then why does it seem possible to respond to them as we respond to reasons, and to lack intention on the basis of them, as we form intentions on the basis of reasons?

Finally, a fourth problem with this defense against my argument is that it violates a key methodological principle: if it quacks like a duck, it's a duck. What I've been arguing, is that the *very same* features which motivate taking the distinction between epistemic and pragmatic reasons for belief and the distinction between normal and toxin-puzzle reasons for intention as being two instances of a single underlying phenomenon, *also* motivate seeing the same underlying contrast among reasons *not* to intend. If the idea that there is something in common between these cases is well-motivated, then that same motivation extends to reasons not to intend. Consequently, being unmoved by the earmarks of a right-kind/wrong-kind contrast in reasons not to intend on the grounds that not intending isn't itself an attitude is a bit like holding that whales can't be mammals, because they're fish – it's true that the object-

given/state-given theory rules out the possibility of right-kind reasons to withhold intention, but that just shows that it was motivated on the basis of an insufficiently broad range of data.

The positive picture that I have been offering contrasts favorably with the primitive sufficiency view on each of these counts. On my picture, we *do* have a positive account of the nature of sufficiency and an account of what makes reasons sufficient: reasons are sufficient just in case they are at least as good as the reasons for each of the alternatives. This picture also offers a simple explanation about *why* the factors that raise the bar on sufficiency do so – they do so because they are reasons for one of the alternatives – withholding intention. Since reasons to intend are sufficient only if they are at least as good as the reasons for each of the alternatives, they need to be better in order to be sufficient, whenever the reasons for any of the alternatives are better. So reasons to withhold intention raise the bar on how good reasons to intend have to be, in order to be sufficient.

As a result, on my picture, the distinction between the ‘right’ and ‘wrong’ kind of factors to raise the bar on how good reasons to intend have to be, in order to be sufficient, is just a special case of the distinction between the ‘right’ and ‘wrong’ kind of reasons, in the first place. Consequently, rather than needing to solve a *new* problem about distinguishing between which factors raise the bar on sufficiency and which do not – which appears to have much of the same structure as the problem with which we started – we only have one problem to solve, about what makes reasons of the right kind. I conclude that my positive picture has greater explanatory power and deeper theoretical attractions than the primitive sufficiency view.

D Consequences

1 *Distinguishing ‘Right-Kind’ State-Given Reasons from ‘Wrong-Kind’ Ones*

In part B I argued that any way of developing the object-given/state-given theory will run afoul of the problem that there are right-kind reasons not to intend (and not to believe), but these reasons can’t be object-given. It should also be pointed out that the reasons on which I based my argument *are*, intuitively, state-given, in the sense that they derive from benefits or costs of being in certain states of mind – having or lacking intentions or beliefs. In the case of my decision about whether to drive to LA, the reason to lack either intention derives from a *benefit* of lacking either outright intention²² – that it allows me to take further information into account in making my decision (similar points go for belief, for the case of the biopsy). Similarly, in the case of my decision about whether to enroll at Old School Eastern or Western Powerhouse, the reason not to form either intention derives from benefit of lacking either intention – that

it allows me to more efficiently take care of my smoke alarm (and the garbage) and allows me to decide later, when I have more time to reflect. So not only are these reasons not object-given, they all seem to be closely connected to *benefits* or *costs* of having or lacking beliefs or intentions – that is, to be *state-given*.

This means that whatever the distinction between the ‘right’ and ‘wrong’ kinds of reasons turns out to be, it is going to have to tell us *which* benefits or costs of being in a certain state of mind are the *right kind* of benefits or costs to generate right-kind reasons, and which are the *wrong kind* of benefits or costs to generate such reasons. The whole idea of the object-given/state-given theory is that we can *avoid* having to distinguish between different kinds of benefits and costs of being in different states of mind, because benefits and costs only give rise to the wrong kind of reason. But if what I’ve been arguing here is on the right track, then only by distinguishing between *kinds* of benefits or costs are we going to have a sufficiently general approach to understanding the distinction between the right and wrong kind of reason.

So far, we’ve seen that the general form of an answer to what distinguishes right-kind reasons from wrong-kind reasons needs to distinguish between right and wrong kinds of *benefit*, or *cost*. But it’s very unpromising to think that there is any once-and-for-all distinction which will do this work, whereby some kinds of benefits are the right kind to ground right-kind reasons for belief, for intention, for fear, for admiration, and so on, and all other kinds of benefits can ground only wrong-kind reasons, whether for belief, intention, fear, admiration, or whatever. In particular, whatever the distinction in kinds of benefits is, it is much more promising to think that there is a different distinction among the benefits of belief than among those of intention. This is, roughly, the grounds on which I claimed in previous work that the right-kind/wrong-kind distinction is relative to what I misleadingly called an *activity* – that is, there are right and wrong kinds of reasons with respect to *belief*, right and wrong kinds of reasons with respect to *intention*, and, if my earlier arguments are on the right track, even right and wrong kinds of reasons with respect to *tying knots*.

So the general form of an account of the distinction between right and wrong kinds of reasons will have to give us some *mapping* from what I’ve called *activities* – believing, intending, fearing, tying knots, and so on – to the kinds of benefits that give rise to the right kind of reason with respect to that activity. In other words, something about the difference between belief and intention will have to tell us what makes the difference between the kinds of benefits that give rise to the right kind of reason for or against belief, and the kinds of benefits that give rise to the right kind of reason for or against intention. In other words, the class of right-kind reasons with respect to any activity will need to depend on the nature of that activity, in some way.

But note that there *is* an intelligible connection between the right kind of reasons for and against intention that we've been considering, and the nature of intention – the kind of attitude that it is. According to a natural hypothesis, intention is an attitude whose point is to close off deliberation, in order to allow us to coordinate and control our own actions across time and make decisions at times at which we have more available cognitive resources.²³ Since closing off deliberation is part of the point and cognitive role of intention, the fact that being undecided about whether to drive to LA allows me to take further information into account is in fact closely related to the role of intention – which contrasts sharply with the fact that I could gain a financial reward by having or lacking an intention, which has nothing in particular to do with what is distinctive of intention at all. The cost of intending to drive into LA only arises if that intention plays its ordinary role of closing off my future deliberation so that I don't take into account the new information provided by my brother's phone call. But the cost of intending when I've been offered money not to arises whether or not that intention plays its ordinary role.²⁴

Similarly, the fact that being undecided about whether to enroll at Old School Eastern or Western Powerhouse enables me to take care of more pressing business right now and decide about my graduate school career later when I have more time to think about it is itself closely connected to the role of intention, of allowing us to efficiently allocate our decision-making cognitive resources over time, so that we can make decisions in 'down time' about what to do at times that we will be busy. Again, this contrasts sharply with 'wrong kind' reasons to intend or lack intention, which derive from benefits that have nothing in particular to do with the distinctive role of intention, and which accrue to it whether or not it plays that role. And even the fact that I need to coordinate now with my wife about who will get the car is closely related to the role for intention, which, as Bratman has described it, is intimately connected to allowing for such coordination of action. This again contrasts sharply with a financial reward for making up my mind, which has no particular connection to the distinctive role of intention.

Just as these observations suggest that our examples of right kind reasons not to intend are closely connected to benefits of not intending that derive from the distinctive role of intention, it is also true that paradigm right kind reasons *to* intend are closely connected with benefits of intention that derive from its distinctive role. After all, since intention's distinctive role is to close off deliberation in order to lead you to act at some later time, any reason *to* do A at that later time will itself yield a corresponding benefit of intending to do A at that time. So even the paradigms of object-given reasons can be easily accounted for as state-given – deriving from the benefits of being in that state – and these benefits are closely connected with the distinctive role of intention – in contrast with wrong-kind reasons.²⁵

Similar points go for belief; for again according to an attractive though somewhat more controversial hypothesis, the distinctive role of all-out belief, as opposed to credence, is to substantially simplify the computational demands on thinking by being a state that effectively closes off uncertainty so that we have something to rely on, in reasoning.²⁶ On this picture, we would expect right-kind reasons for and against belief to be closely associated with benefits or costs of believing that are closely connected to this distinctive role of closing off uncertainty so that we have something to rely on, in reasoning. And again, this is highly plausible – the benefit of waiting to close off uncertainty in the biopsy case is that it allows us to take into account further evidence. And in general, evidence that p makes it more likely that in believing that p , you will be reasoning from the truth – so insofar as there are benefits to relying on the truth in your reasoning, evidence will itself be closely connected to benefits which derive from the distinctive role of belief.

Note that in this section I have *not* been trying to outline a theory about what ultimately distinguishes the right kind of reason from the wrong kind. I have tried to offer a more concrete positive theory in previous work, but I am now agnostic about its details. What I *have* been arguing is that however this distinction works, it must work by somehow associating each attitude (or ‘activity’, more generally) with a distinctive kind of benefit or cost which generates the right kind of reason with respect to this attitude (or activity).

I have also been pointing out that the right kinds of reason for and against intention and belief are clearly associated with benefits or costs of intention and belief that are much more closely connected to the nature of intention and belief than the benefits and costs are in cases involving wrong kind reasons. I haven’t tried to articulate exactly *how* these benefits and costs are associated with intention and belief, but I take this to be some kind of evidence that some account with this general form will be able to successfully distinguish right kind reasons from wrong kind ones. Here all I aim for, is to illustrate the more promising virtues of this general strategy, in contrast to that of the object-given/state-given theory.

2 *High-Handed Arguments*

I close by briefly mentioning two important arguments which rely on the object-given/state-given theory, in order to illustrate some of the reasons why it is important that this theory is false. Each of these arguments uses the object-given/state-given theory in order to argue that certain theories *can’t* be right. I call these arguments ‘*high-handed*’ because rather than engage with the details of the particular views that they address, they presume on the strength of the obviousness of the object-given/state-given theory to be able to discern that no theory in some broad domain could even *possibly* be right.

The first such high-handed argument is inspired by a loose reading of Derek Parfit's recent arguments against desire-based theories of reasons.²⁷ According to desire-based theories, every reason is explained by some desire of the agent for whom it is a reason. According to this objection, however, such theories have the consequence that it is impossible for a desire to be rational or irrational – so since it *is* possible for desires to be rational or irrational, all such theories are false. Though I explained in *Slaves of the Passions* how it *is* possible for a desire-based theory to accommodate and explain *reasons* for and against having desires, this is not what this objection denies.²⁸ Rather, the objection turns on the point that desire-based theories can only accommodate *state-based* reasons to have desires – for a desire-based theory can explain a reason to have a desire only if *having that desire* would serve to promote the object of some other desire, and that is a paradigmatically state-given reason – a benefit of being in the *state* of having that desire.²⁹ Hence, according to this objection, desire-based theories can't explain the right kind of reason for or against desire, and hence can't allow that desires can be rational or irrational.

Clearly, this objection trades on the object-given/state-given theory. If state-given reasons can be of the right kind, as I have been arguing in this paper, then that opens up the possibility that a desire-based theory of reasons can accommodate and explain the *right kind* of reasons for and against desire – and for and against intention and belief, for that matter. In order to evaluate whether this possibility is actual, we again need to get past the high-handedness and evaluate in detail what makes reasons for and against desire of the right kind. And if my arguments in this paper are correct, that means that we need to look more carefully at the nature of desire.

The second high-handed argument I want to mention comes from what I think is actually a very common reaction to views proposing what has come to be known as 'pragmatic encroachment' in epistemology. Recently several authors, including John Hawthorne, Jason Stanley, and Jeremy Fantl and Matthew McGrath, have endorsed the thesis that whether someone knows *p* can depend on – in addition to whether *p* is true, her evidence for *p*, the reliability of her environment and her cognitive faculties, and such other truth-related factors – factors like how much is at stake for her over whether *p* is true.³⁰ In a standard sort of example, suppose that you are driving past the bank on your way home from work and have your paycheck to deposit. But the lines are long, and you'd prefer to come at another time with shorter lines, if possible. Fortunately, you believe that the bank will be open Saturday morning. According to the pragmatic encroachers, in at least some cases like this, you need more evidence that the bank will be open on Saturday, if your paycheck needs to be deposited by Sunday in order to cover your bills, than if it would be just as good for you if the check were deposited on Monday. Thus, whether you know can depend on the *stakes* – a non-evidential, non-truth-related, factor.

Though these authors have offered a number of reasons for accepting this view, it has seemed to many like the kind of thing that *can't* be true. After all, the intuitive thought goes, knowledge can only depend on the evidence, or at least only on truth-related factors, because only evidence can bear on *epistemic* rationality, as opposed to the broadly pragmatic sort of rationality that endorses believing in accordance with Pascalian considerations.³¹ This line of thought plausibly implicates the object-given/state-given theory, as it relies on the assumption that only evidence can affect distinctively epistemic rationality. But in this paper I've been giving reasons to doubt the object-given/state-given theory, and in particular, to suspect that there are non-evidential reasons which bear on distinctively *epistemic* rationality – including the fact that further, decisive, evidence will soon be forthcoming, as in the biopsy case.

Indeed, high-stakes cases can also be directly used to motivate the rejection of the object-given/state-given theory. For example, it is relatively easy to withhold belief about whether the bank will be open on Saturday, on the grounds that so much is at stake, if you are wrong (for example, that your mortgage payment is due on Sunday and you are nearing foreclosure). Knowing the stakes are high is exactly the kind of thing to make you hold out for more evidence before making up your mind – again contrasting sharply with monetary offers. It is controversial, of course, whether knowing the stakes are high makes it *epistemically rational* to withhold belief, but the main reason this is controversial is that people are wedded to the idea that the only right-kind reasons for and against belief are evidence, which is precisely what is here in question.

Moreover, in keeping with the line of thought pursued in section DI, the costs of being wrong in such cases are closely connected to the role of belief in closing off uncertainty so that you have something to rely on in your reasoning. After all, the relevant cost of being wrong in your belief that the bank will be open on Saturday is that because you *count on it to be open*, you go home on Friday afternoon and come back Saturday morning to deposit your paycheck, but since you were wrong, the bank is closed, so your paycheck doesn't get deposited and you miss your deadline and the bank forecloses on your house. These costs of being wrong are intimately tied to the fact that your belief is doing its normal job – giving you something to rely on in your reasoning. Again, this contrasts sharply with the wrong-kind cases.³²

In this section I've briefly mentioned two kinds of high-handed argument that seem to rely on versions of the object-given/state-given theory. I don't introduce them in order to resolve the issues at stake, but to illustrate why it is so important whether the object-given/state-given theory is true. If it *is* true, then some genuinely exciting conclusions follow. If it's not, well, much more work is needed, including more work on the nature of desire and belief. I conclude that much more work is needed.

3 *Conclusion*

In this paper, I've sought to leverage the reasons that we have *against* intention and belief, in order to explain why the object-given/state-given theory is on the wrong track, in trying to characterize what the difference is between the 'right' and 'wrong' kinds of reasons – the contrast that we see between epistemic and pragmatic reasons for belief, and between normal and toxin-puzzle reasons for intention. To do so, I've focused on what I hope is a helpful way of framing the problem about the distinction between the right and wrong kind of reasons, by framing it in terms of earmarks of the distinction.

I've sketched the most general structure of an alternative, I think much more promising, strategy for understanding the difference between the right and wrong kinds of reasons – a structure which, like the basic intuition behind the object-given/state-given theory, can be filled out in a wide variety of different ways, whose differences I'm neutral about, here. And finally, I've illustrated some of the important issues in moral philosophy and epistemology which are directly affected by whether the object-given/state-given theory is on the right track or not.

Notes

¹ In fact there are two different ways in which the distinction between the ‘right’ and ‘wrong’ kind of reasons entered the literature; one way, which will not be a primary focus of interest in this paper, is through the dialectic surrounding ‘fitting attitudes’ and ‘buck-passing’ accounts of value. See, in particular, Wlodek Rabinowicz and Toni Ronnow-Rasmussen, “The Strike of the Demon: On Fitting Pro-Attitudes and Value,” *Ethics* 114 (2004): 391-423, and Mark Schroeder, “Value and the Right Kind of Reason,” *Oxford Studies in Metaethics* 5 (2010): 25-56.

² For discussion, see especially Bernard Williams, “Deciding to Believe,” in his *Problems of the Self* (Cambridge: Cambridge University Press, 1973): 148; Jonathan Bennett, “Why is Belief Involuntary?” *Analysis* 50 (1990): 93; and Pamela Hieronymi, “Controlling Attitudes,” *Pacific Philosophical Quarterly* 87 (2006): 45-74.

³ The fourth, ‘flavor’ earmark is arguably not a proper earmark in its own right, but just the report of the intuitive naturalness of classifying various putatively wrong-kind reasons together. Its exact status won’t matter for my purposes, here.

⁴ Gregory Kavka, “The Toxin Puzzle,” *Analysis* 43 (1983): 33-36.

⁵ See especially Thomas Pink, *The Psychology of Freedom* (Cambridge: Cambridge University Press, 1996), David Owens, *Reason Without Freedom* (New York: Routledge, 2000), and Pamela Hieronymi, “The Wrong Kind of Reason,” *Journal of Philosophy* 102 (2005): 437-457.

⁶ See especially Justin D’Arms and Daniel Jacobson, “Sentiment and Value,” *Ethics* 110 (2000): 722-748 and Justin D’Arms and Daniel Jacobson, “The Moralistic Fallacy,” *Philosophy and Phenomenological Research* 61 (2000): 65-90.

⁷ See Rabinowicz and Ronnow-Rasmussen, “The Strike of the Demon” and Schroeder, “Value and the Right Kind of Reason” for some of the reasons this question is interesting that aren’t discussed in this paper.

⁸ Compare especially John Skorupski, “Buck-Passing About Goodness,” in Ronnow-Rasmussen, Petersson, Jesefsson, and Egonsson, eds., *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, www.fil.lu.se/hommageawlodek, but also Derek Parfit, *On What Matters*, volumes I and 2 (Oxford: Oxford University Press, 2011) and Jonathan Way, “Transmission and the Wrong Kind of Reason,” forthcoming in *Ethics*.

⁹ See particularly Derek Parfit, "Rationality and Reasons," In Dan Egonsson et al, eds., *Exploring Practical Philosophy* (Burlington, VT: Ashgate, 2001), 17-39; Christian Piller, "Normative Practical Reasoning," *Proceedings of the Aristotelian Society* Suppl. Vol. 25 (2001): 195-216; Christian Piller, "Content-Related and Attitude-Related Reasons for Preferences," *Philosophy* Suppl. Vol 59 (2006): 155-181; Jonas Olson, "Buck-Passing and the Wrong Kind of Reasons," *The Philosophical Quarterly* 54 (2004): 295-300; and Philip Stratton-Lake, "How to Deal with Evil Demons: Comment on Rabinowicz and Ronnow-Rasmussen," *Ethics* 115 (2005): 778-798. The view offered by Pamela Hieronymi, "The Wrong Kind of Reasons," differs in important ways from most of the views in this family, but shares the commitments that will make trouble in what follows, so it is also one of my targets.

¹⁰ This isn't exactly right; one important complication is that 'you want to' mandates a *de se* perspective on what is wanted, whereas 'want yourself to' does not. This complication is not important for the point in the main text.

¹¹ It is worth noting that although the view in Hieronymi, "The Wrong Kind of Reason," will be among the targets of my arguments in part B, it does not face this particular problem.

¹² Some of the complications in formulating this theory in an adequate way are pursued, in particular, in Rabinowicz and Ronnow-Rasmussen, "The Strike of the Demon," Olson, "Buck-Passing," Stratton-Lake, "How to Deal with Evil Demons," Wlodek Rabinowicz and Toni Ronnow-Rasmussen, "Buck-Passing and the Right Kind of Reasons," *The Philosophical Quarterly* 56 (2006): 114-120, Sven Danielson and Jonas Olson, "Brentano and the Buck-Passers," *Mind* 116 (2007): 511-522, Gerald Lang, "The Right Kind of Solution to the Wrong Kind of Reason Problem," *Utilitas* 20 (2008): 472-489, and Hieronymi, "The Wrong Kind of Reason."

¹³ I'll further defend this characterization of the commitments of the object-given state-given theory in section B3.

¹⁴ Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press, 2007), Schroeder, "Value and the Right Kind of Reason".

¹⁵ Sometimes lacking either intention is referred to, following similar terminology from epistemology for the case of belief, as *withholding*. But some philosophers would prefer to distinguish withholding intention from simply lacking intention. So I will avoid using this term. I will return in section CI to consider whether this distinction is of importance.

¹⁶ Thanks to Don Hubin for pressing me on this.

¹⁷ Careful readers will observe both that this falls short of an endorsement of **intend okay**, and that it raises the question as to whether I was justified in attributing a commitment to **intend object** to proponents of the object-given/state-given theory. I'll return to both of these issues in section B3.

¹⁸ As before, sometimes 'withholding' is used as a term for the state of lacking either belief. But some philosophers would prefer to distinguish withholding from mere lack of belief, reserving 'withholding' for intentional reservation of judgment. Again, I'll come back in section CI to whether this distinction is of importance.

¹⁹ There is one wrinkle, when it comes to belief, that does not arise for the case of intention. And that concerns the relationship between *binary* belief and *credence*, or degree-theoretic belief. The claim I've just made – that in addition to believing p and believing $\sim p$, there is a third option of lacking either belief – is a claim about *binary* belief. But in addition to the pretheoretic notion of binary belief, many philosophers hold that we also have various degrees of confidence which can be represented by something like a probability function, and which are commonly known as *credences*. According to some philosophers, the existence of *credences* obviates the need to think that there is any such thing as binary belief at all, and according to others – *credal reductivists* – talk about binary belief can be reduced to talk about credences.

I will assume here, however, in accordance with much work in traditional epistemology, that these views are mistaken. I believe that it is safe to assume that there is a real psychological difference between being unopinionated with respect to p and having made up one's mind with respect to p that does not track any particular degree of credence. In short, something happens when you make up your mind, but you can make up your mind without being fully confident, and you can be highly confident without having made up your mind. So binary belief – the state of mind that involves making up your mind about p – is its own psychological state and does not reduce to credence. At any rate, I will assume so in what follows; for argument in defense and elaboration of this view see especially Jacob Ross and Mark Schroeder, "Belief, Credence, and Pragmatic Encroachment," forthcoming in *Philosophy and Phenomenological Research*.

²⁰ Again, careful readers will note that this falls short of an endorsement of **believe okay**.

²¹ Note that this is the reverse of my earlier case, in which in the absence of definitive information about where my brother would be, the rational choice was to stay home.

²² Note that this benefit of lacking either outright intention may also be shared by someone who has a *conditional* intention. So it is not, strictly speaking, a benefit of being *undecided*. Thanks to Don Hubin for helping me get clear on this point.

²³ See, in particular, Michael Bratman, *Intention, Plans, and Practical Reason* (CSLI: Stanford, 1987), and Pink, *The Psychology of Freedom*. Note that though intention closes off deliberation with respect to some particular question, it can leave much still open – in forming an intention to drive to LA, for example, I can leave open which route to take.

²⁴ Clearly there will be subtleties involved in making this idea precise. For example, if I am instead offered money for not having an intention that plays its ordinary cognitive role, that still looks like a wrong-kind reason to lack the intention.

²⁵ Compare especially Pink, *The Psychology of Freedom* on the role of intentions as means.

²⁶ Compare Gilbert Harman, *Change in View* (Cambridge, MA: MIT Press, 1988); but see especially Ross and Schroeder, “Belief, Credence, and Pragmatic Encroachment” for an elaboration and defense of the cognitive role of binary belief.

²⁷ Parfit, *On What Matters*. This argument actually does not appear in the final print edition of *On What Matters* but was offered to me by Jake Ross as inspired by Parfit’s reasoning. Whether or not the argument is Parfit’s, I think the objection is important and I’m indebted to Ross for helping me to appreciate its force.

²⁸ Schroeder, *Slaves of the Passions*, Chapter 10.

²⁹ See especially chapter 10 of *Slaves of the Passions*.

³⁰ John Hawthorne, *Knowledge and Lotteries* (Oxford: Oxford University Press, 2004), Jason Stanley, *Knowledge and Practical Interests* (Oxford: Oxford University Press, 2004), Jeremy Fantl and Matthew McGrath, “Evidence, Pragmatics, and Justification,” *Philosophical Review* III (2002): 67-94, Jeremy Fantl and Matthew McGrath, *Knowledge in an Uncertain World* (Oxford: Oxford University Press, 2010).

³¹ Compare Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), chapter 1. See especially Mark Schroeder, “Stakes, Withholding, and the Subject-Sensitivity of Knowledge,” forthcoming in *Philosophical Studies*, for discussion.

³² Compare Schroeder, “Stakes, Withholding” and Ross and Schroeder, “Belief, Credence, and Pragmatic Encroachment.”