What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)

Errol Lord
University of Pennsylvania
erlord@phil.upenn.edu

It is a truism that we ought to be rational. Despite this (or because of it), it has become popular to think that it is not the case that we ought to be rational. In this paper I argue for a view about rationality—the view that what one is rationally required to do is determined by the normative reasons one possesses—by showing that it can vindicate that one ought to be rational. I do this by showing that it is independently very plausible that what one ought to do is determined by the normative reasons one possesses. Thus, the paper also makes a contribution tos the debate about the nature of our obligations.

1. Introduction

It has been notoriously difficult to vindicate the natural thought that moral considerations necessarily bear on what one ought to do. That is to say, it is unclear whether the facts that determine what one is morally required to do necessarily affect what one ought to do, full stop. The requirements of morality thus face a sort of sceptical challenge. The challenge is to explain the *deontic significance* of morality. Deontic significance comes in at least two forms. On the one hand, it might be that we always ought to do the thing that morality requires. On the other hand, it might merely be the case that we always have reason to do what morality requires. Thus, in order to meet the sceptical challenge, one must explain why we always ought to do or have reason to do what morality requires.

Despite this scepticism about morality, other domains have traditionally been seen to be on steadier deontic footing. A paradigm case is *rationality*. Rationality, it is often thought, is clearly deontically significant. The considerations that bear on what is rational seem to necessarily affect what one ought to do, full stop. In fact, something stronger seems to be true, namely, that you always ought to be

Mind, Vol. 126 . 504 . October 2017
doi:10.1093/mind/fzw023 Advance Access publication 19 January 2017

rational. It does not seem like much of a challenge at all to explain why we always have most reason (and thus some reason) to do what rationality requires. It is a striking fact that many of the classic positions—both sceptical and non-sceptical—concerning the deontic significance of morality assume the deontic significance of rationality.[1] This is evidence enough that rationality is usually seen to be on steady deontic ground.

Despite this (or perhaps because of it), it has become increasingly popular in the literature on practical reason to think that it is not necessarily the case that you should be rational.[2,3] And given some assumptions about rationality and about what determines what you should do, this actually seems plausible. Moreover, those assumptions about rationality and about what determines what you should do are thought to be supported by powerful arguments. Because of these rather surprising facts, scepticism about the deontic significance of rationality is alive and kicking. Back to this in a moment.

Another corner of normative philosophy is concerned with what determines what one ought to do, full stop. This debate is about what the best theory is of what we might call the *ought of deliberation*. This is the 'ought' that figures in the central deliberative question, 'What ought I to do?'[4] It's commonly assumed that the answer to the central deliberative question is the thing that you ought to do, full stop—it *settles (correct) deliberation*. What you ought to do and believe full stop is, I will say, what you ought to do and believe. The debate I'm interested in here is about which facts *determine* what you ought to do (and believe). The relevant question is this: do all of the facts determine what you ought to do or do only some of them?

Consider an example.[5] Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling art historian/dealer, doesn't have the money. It

[1] This breaks down roughly into Humean sceptics and Kantian rationalists. On the sceptical side, see—to name just two—Foot (1972) and Harman (1985). On the rationalist side, see Smith (1994) and Korsgaard (1996).

[2] I will follow the standard practice of treating 'should' as a synonym for 'ought'.

[3] The most prominent sceptic is Niko Kolodny (see especially Kolodny 2005 and 2007). Despite not being a sceptic, John Broome has also done a lot of work arguing against the claim that rationality is deontically significant (this work is brought together nicely in Broome (2013).

[4] In setting things up this way I'm following (among others) Broome (2013), Kiesewetter (2011), Ross (2012), and Kolodny and MacFarlane (2010).

[5] I consider this example along with many of the cases in section 3 in Lord (2015).

looks as if his mother is going to die. She will, however, be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go to see her. As it happens, a pawnshop owner in Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. The rub, of course, is that he has no idea that this pawnshop even exists, much less that such a deal is to be had there. Interesting question: what ought Jack to do with his day? Go to California or go to Queens?[6]

If you think that all of the facts are eligible to determine what one ought to do, then you are committed to thinking that the place for Jack to go is Queens. If you don't like that answer, then you should think that only some of the facts are eligible to determine what one ought to do. The trick is to explain which facts are relevant.

There is a tight connection between the correct resolutions to these two debates—the debate about whether you ought to be rational and the debate about what determines what you ought to do. I will argue in this paper for a neglected theory of rationality—the *Reasons Responsiveness* account—by showing that it vindicates the deontic significance of rationality. According to the Reasons Responsiveness account, rationality consists in correctly responding to the normative reasons you possess, where the reasons you possess are the facts that count in favour of acts and attitudes that are within your ken. The Reasons Responsiveness account holds that what you are rationally required to do is determined by the normative reasons you possess. If this view is correct, I will argue, then what you ought to do *just is* what you are rationally required to do. In order to show this, I will argue for a view about what determines what you ought to do. I will argue that there are strong reasons to think that what determines what you ought to do are the reasons you possess. The basic idea behind my argument for this is that the reasons that obligate have to be potentially action-guiding in a certain sense—we have to be able to act *for* the reasons that obligate. I will argue that in order to act for a reason in the appropriate way, one has to possess the reason. It is a few short steps from here to the claim that the reasons one possesses determine what one ought to do.

[6] Some think that there isn't really any good question here. This is because, claim these theorists, all we need to do to settle the dispute is distinguish between a subjective and objective sense of 'ought'. I disagree that this is enough to solve the problem, as do many others. I spell out why I disagree in §4.1.

The paper has two main upshots. First, it provides a defence of a specific way of analysing rationality in terms of normative reasons. While this kind of view seems to many to be very intuitive, it has fallen on hard times in the metaethics literature. As we'll see, the direction that the literature about rationality in metaethics has taken has led to a sort of quagmire. The main upshot of this paper is that the quagmire can be avoided by returning to theorizing about rationality in terms of reasons. Further, I show this in a novel way—namely, by showing that Reasons Responsiveness fits nicely with a plausible view about what determines what you ought to do. The second main upshot is that the paper provides a new argument for a view about what determines what we ought to do. Thus, the paper not only makes headway in the stalled literature about the deontic significance of rationality, it also helps to advance an important view in the debate about what determines our obligations.

The paper has the following structure. I begin by canvassing the recent literature in metaethics about rationality and its deontic significance. The upshot will be that the two most discussed views have serious problems vindicating the deontic significance of rationality. I will use this section to initially motivate the Reasons Responsiveness account. The Reasons Responsiveness account immediately entails that we always possess reasons to be rational. Unfortunately, it doesn't immediately follow from the Reasons Responsiveness view that one always ought to be rational. In order to settle this question, we must investigate the second debate. So in the second half of the paper we will focus on this. I will argue that it is independently very plausible that what you ought to do is determined by the reasons you possess. Given that the reasons you possess determine what you are rationally required to do, according to the Reasons Responsiveness view, the view vindicates the full deontic significance of rationality. This is a powerful reason to accept the view.

2. Why be rational?

In the middle of the last decade, a wave of philosophers started to question the deontic significance of rationality.[7] Unfortunately, what exactly is being questioned is itself a matter of debate. Because of this, I'll start by explicating what I take to be the core of the challenge. I will

[7] See Kolodny (2005, 2007, 2008a, 2008b), Broome (2005a, 2005b, 2008b, 2013), Raz (2005), Southwood (2008), and Way (2010a, 2010b).

then introduce the main views discussed in the literature, and show how they have a hard time meeting the challenge. This will be by way of motivating my positive account, which will be introduced in §3.

2.1 Deontic significance, what

The easiest way into the question being asked is to think about the most obvious way in which rationality could turn out to be deontically significant. Rationality would obviously be deontically significant if one always *should* be rational. This is not simply to say that one always *rationally should* be rational—that one is always required by rationality to do the things rationality requires. It's not open to dispute whether one always rationally should be rational. Rather, the question is whether one always should be rational, full stop.

The most prominent way to think about the question in the recent literature is to imagine that there are different systems of requirements—for example, morality, prudence, etiquette and rationality.[8] Each of these different systems will have particular *sources*, which will consists in some considerations that bear on what one is required to do by that system. A paradigmatic system in this sense is the law. A particular set of legal requirements has its source in a legislature and a system of precedent. This source gives rise to a set of requirements—the legal requirements. The recent literature about rationality has thought of rationality as analogous to the law in this way. There will be some source that gives rise to a set of requirements (which considerations are in the source for the system of rational requirements will be discussed in the next subsection).

Importantly, on this view, there is no conceptual connection between the different systems of requirements and normative reasons or the all-things-considered notion of obligation. In the case of the law, this is plausible. It is very plausible that at least some legal systems are divorced from the actual normative facts. Thus, it seems as though it will be an open question whether there are normative reasons to do what the law requires, and it will be an open question whether one ought to do what the law requires. Since rationality functions like the law, on this view, similar questions can be asked about rationality. Once we determine what rationality requires, we can ask whether there are normative reasons to do the things it requires, and we can ask

[8] This is the most prominent way of thinking of things because this is the way John Broome thinks of the issues; see, for example, Broome (2007b, 2013). Ultimately, I think that this way of thinking of things is misleading, but it is a good way to cotton on to what's going on.

whether one ought to do the things it requires. This is to ask if the rational requirements are deontically significant.

Broome and others have focused on two particular questions regarding rationality. The first question is whether rationality is a system that is such that if it requires one to ϕ, then one thereby ought to ϕ. The second question is whether rationality is such that if it requires one to ϕ, then one thereby has a reason to ϕ.

Despite the fact that the literature has focused on these rather narrow questions, I don't think this is the best way to see things. These questions are *too narrow*. This is because they assume that rationality is deontically significant only if *the fact that rationality requires ϕ-ing* provides either a full-stop requirement to ϕ or a reason to ϕ. I agree that these are two ways that rationality could turn out to be deontically significant. But we can also ask the broader questions of whether it's true that one always ought to do what rationality requires or whether it's true that one always has a reason to do what rationality requires. It might be true that one always ought to do what rationality requires or has reason to do what rationality requires *even if* those facts aren't grounded in the fact that rationality requires one to do the thing in question. Let's focus, then, on the broader questions. To ease discussion, let's say that rationality is *strongly (deontically) significant* if it turns out that one always ought to do what rationality requires. And let's say that rationality is *weakly (deontically) significant* if it turns out that one always has reason to do what rationality requires.[9]

2.2 Coherence, rationality, deontic significance

This section is dedicated to explicating how rationality is conceived of in the recent literature about the deontic status of rationality. The literature has focused on two different views. Both views are motivated by the thought that rationality has a tight connection with being coherent in certain ways. Common examples are means–end coherence, belief consistency, and following your conscience. To save words, I will focus mostly on means–end coherence. Both of the views that the literature has focused on think that one is rational when one is means–end coherent and that one is irrational when one is means–end incoherent.

[9] Note that this way of thinking about deontic significance is still compatible with thinking that each set of requirements has a source in Broome's sense. This compatibility is a dialectical advantage; that said, in the conclusion I will explain how my view dispenses with Broomian sources.

The two views offer different explanations of why one is rational when one is coherent and irrational when one is incoherent. According to the first view, one is irrational when means–end incoherent because one violates *Means–End N*:

> ***Means–End N***: If *A* intends to ϕ and believes that in order to ϕ they must intend to ψ, then *A* is rationally required to intend to ψ.

Since 'rationally required' just scopes over the consequent of *Means–End N*, we'll call it a narrow-scope requirement. Narrow-scopers think that one is irrational when means–end incoherent because one lacks an intention rationality requires one to have.

Usually the narrow-scope view is introduced merely as a foil for what has come to be the dominant view (more on why it has become the foil below). This view holds that one is irrational when one is means–end incoherent because one violates *Means–End W:*

> ***Means–End W:*** *A* is rationally required to [intend to ψ if they intend to ϕ and believe that they must intend to ψ in order to ϕ.]

'Rationally required' takes wide scope over the conditional in *Means–End W*. Thus, it is a wide-scope requirement. Wide-scopers think one is irrational when means–end incoherent because a conditional one is required to make true is false.

The wide-scope view, being the dominant view, has been the main target of the sceptical arguments.[10] Its failure to account for the deontic significance of rationality is all the more surprising because it has been influentially motivated by the thought that only it can account for why rationality is deontically significant.[11] The argument that often

[10] Kolodny (2005) brought scepticism to the fore. Others who raise the sceptical challenge include Broome (2013), Raz (2005) and Way (2010b, 2012).

[11] The thought goes back at least to Hill (1973), Greenspan (1975) and Darwall (1983). It became widely accepted because of the work of John Broome, especially Broome (1999). In recent years, people have become much less confident about (3) (Broome still endorses the claim that rationality is deontically significant in Broome 2013, even though he admits he doesn't have a persuasive argument for this). This is because of the arguments given by theorists like Kolodny and Broome. Thus, many people do not now wholly endorse this argument. But I think it's still fair to say that the consensus view is that the wide-scope requirements are the only requirements with any hope of being deontically significant. This is because people are still by and large convinced by the argument for (2). Further, the argument that I'm about to present in favour of (2) is taken by most to be fatal to the narrow-scope view independently of the debate about the deontic significance of rationality.

motivates the wide-scope view goes something like this (cf. Schroeder 2004, 2005a):[12]

(1) Either the wide-sscope view is true or the narrow-scope view is true.

(2) If the narrow-scope view is true, then rationality is not deontically significant.

(3) Rationality is deontically significant.

(4) Thus, the narrow-scope view is false.

(C) Therefore, the wide-scope view is true.

Let's start with (2). If the narrow-scope view were true and rationality were deontically significant, then rationality would give rise to objectionable bootstrapping. To see this, suppose John intends to eat a planet made of Stilton cheese and believes that in order to do that he needs to intend to build a spaceship. If *Means–End N* is true, it thereby follows that rationality requires John to intend to build a spaceship. This is in itself implausible. But it gets even worse if you think that rationality is deontically significant. If it's strongly significant, then it would follow that John ought, full stop, to intend to build a spaceship. This is certainly false. If it's weakly significant, then it would follow that John has a reason to intend to build a spaceship. This isn't as obviously false as the previous claim, but it still strikes nearly everyone as very implausible.[13]

Means–End W doesn't have these problems. This is because John can comply with *Means–End W* without intending to build a

[12] As an anonymous referee has pointed out to me, there is another route to scepticism about the deontic significance of rationality that does not go through the idea that rationality is tied to coherence. Perhaps the paradigm here is Parfit (1997). Parfit holds that rationality is a function of the apparent reasons, and holds (rightly) that apparent reasons aren't deontically significant. One way of seeing the history is that the difference between early Broome and early Parfit is that they both recognized that (2) is true, but had different reactions to it. Broome wanted rationality to be deontically significant, so he went for the wide-scope view. Parfit didn't mind denying the deontic significance of rationality, so took the truth of (2) as not unwelcome. An important point for my project, though, is that Parfit never really argues for any of this; he just presents it as true. It's also interesting to note that further refinements to Parfit's idea (especially in Schroeder 2009b) try to show that the apparent/subjective reasons view can account for the coherence data.

[13] Schroeder (2004, 2005b) argued that the narrow-scope requirements are weakly deontically significant. He has since given this up: see Schroeder (2009b) for his retraction of the earlier view. More recently, Smith (2016) has argued that intentions provide reasons to take the means.

spaceship. He can comply by giving up his end or his means–end belief. Thus, it doesn't follow from *Means–End W* that John is rationally required to intend to build a spaceship when he intends to eat a planet made of Stilton and believes that he must intend to build a spaceship in order to eat a planet made of Stilton. Thus, there is no bootstrapping even if rationality is strongly or weakly significant.[14]

Despite the wide-scopers' (apparent) ability to avoid bootstrapping, it has become increasingly unclear whether the wide-scope requirements are deontically significant in either the strong or the weak senses. Before seeing why, it's very important to further clarify the question being asked about the wide-scope view.

Suppose I believe I am writing and disbelieve I am writing. It's plausible that I'm irrational. The wide-scoper thinks I'm irrational because I'm rationally required to [not disbelieve I am writing if I believe I am writing]. When asking whether this requirement is strongly significant, we are asking whether I ought to not-disbelieve-I'm-writing-or-not-believe-I'm-writing. We are not merely asking whether I ought to be in some state other than the incoherent one. It's plausible that if I have incoherent beliefs, then I will necessarily be violating *some* requirement. In this case, either my evidence will sufficiently support the claim that I'm writing or it will sufficiently support the claim that I'm not or neither claim will be sufficiently supported. My evidence will never sufficiently support both claims at the same time. Given the plausible assumption that one ought not to believe p when one lacks sufficient evidence for p, it follows that whenever I both believe and disbelieve I'm writing right now, it will be the case that I ought not to have (at least) one of those beliefs. In this case, it's the belief that I am not writing right now.

So we can stipulate that I ought not to disbelieve that I'm writing right now. And it's true that if I comply with that requirement, I will be coherent. It is very important that the literature in question assumes that *this is not to say* that I ought to comply with the wide-scope requirement.[15] In other words, we are assuming there is a gap between being required to do something that guarantees I'll be coherent and being required to be coherent. It's quite plausible that it doesn't

[14] Some argue that the wide-scope view does have a bootstrapping problem in cases where it's impossible in some way to comply in certain ways. See, for example, Greenspan (1975) and Schroeder (2009b).

[15] Participants in the literature haven't been clear on this point. The only places I know of where this is explicitly mentioned are Lord (forthcoming b), Kolodny (2007), Way (2013, n. 15) and Schroeder (2015a).

immediately follow from the fact that you are required to do something that guarantees that *p* obtains that you are required to see to it that *p* obtains. To take a famous example, if I comply with my requirement to post the letter, I will guarantee it is true that *I post the letter or I burn the letter*. It is counter-intuitive that it *follows* that I'm required to post the letter or burn the letter.[16] So it goes for coherence as well. It doesn't follow from the fact that I ought to have some attitude that will guarantee coherence that I ought to be coherent.[17]

Once this is made clear, the problem becomes acute. What would make it the case that we ought to always comply with the wide-scope requirements? Remember, it *can't* merely be that we are always required to do things that guarantee we'll be coherent. In order to show that the wide-scope requirements are strongly significant, we'd have to show that there are reasons over and above the reasons we have for individual attitudes that make it the case that we ought to be coherent. But it's completely mysterious what reasons there could be to always make it the case that we ought to be coherent.

To bring out the mystery, let's look at the three strategies most often pursued in the literature.[18] The first strategy is that you always instrumentally ought to be coherent. According to this strategy, it's always the case that by being coherent you are doing something else that you ought to do. Although it's very plausible that you are always doing something else you ought to do when you are coherent *in certain ways*, it's not plausible that you are always doing something else you ought to do by being coherent in any way. To use our earlier example, it's stipulated that I am doing something else I ought to do when I fail to disbelieve I am writing right now. But it seems very implausible that I am doing something else I ought to be doing if I instead drop my belief I am writing right now and continue to disbelieve I am writing right now. This is a way to do what the wide-scope requirement requires. Thus, it's not

[16] This, of course, is Ross's Paradox (see Ross 1941). For a nice discussion in this context, see Broome (2007b).

[17] Similar points apply to vindicating the weak significance of the wide-scope requirements. We can stipulate that I have reason to not disbelieve that I'm writing right now. And thus I have reason to do something that guarantees I'll be coherent. This is not yet to say that I have reason to be coherent. I have a reason to not disbelieve I'm writing. This is not necessarily a reason to not-disbelieve-I'm-writing-or-not-believe-I'm-writing. When we're asking whether the wide-scope requirements are weakly significant, we are asking whether there are always reasons to be like that, not merely whether there are always reasons to be in states that guarantee that you'll be like that.

[18] Both Broome (2013) and Kolodny (2005, 2007) pursue these strategies at length. See also Way (2010a, 2010b).

generally true that I do something else I ought to do by complying with the wide-scope requirement.

The second strategy claims that one ought to comply with the wide-scope requirements because doing so is constitutive of some activity one is engaged in. Candidate activities include believing, intending and reasoning.[19] The underlying idea is that if one is a creature whose attitudes never cohere in the ways mandated by the wide-scope requirements, then one is not a creature with beliefs and intentions or a creature who reasons. While this might be true, it simply does not follow from this that one always ought to comply with the wide-scope requirements. This is because, *inter alia*, it doesn't follow from this that one ceases being a creature with beliefs and intentions or a creature who reasons *every time* one is incoherent. Perhaps a requirement to *sometimes* be coherent follows from these supposed agential facts, but this is a far cry from a full vindication of the wide-scope requirements.

The final strategy holds that coherence is intrinsically very good, so good that you always ought to be coherent.[20] This is very implausible. It is not that implausible to think that coherence is an intrinsic good, but it does seem very implausible that it is so good that its goodness always makes it the case that one ought to be coherent. This is what would need to obtain for coherence's intrinsic goodness to vindicate the strong significance of the wide-scope requirements.

None of the three strategies seem promising when it comes to vindicating the strong significance of the wide-scope requirements. Do any of them hold promise for vindicating rationality's weak deontic significance? I think the same types of argument can be given against the analogous strategies for vindicating the weak significance of the wide-scope requirements. The instrumental strategy holds that one has a reason to be coherent because by being coherent one will always be doing something else one has reason to do. But it seems possible that one can comply with the wide-scope requirements in ways that are such that one isn't doing something else one has reason to do by complying in those ways. The constitutivist strategy holds that some agential facts ground a reason to be coherent. Again, it's not clear which agential facts could do this. Some agential facts

[19] See Bratman (2009a, 2009b) and Buss (MS) for a defence of the first two, and Hussain (MS) for a defence of the reasoning account.

[20] As far as I know, this position has never been endorsed. It is discussed and argued against by Kolodny and Way in the papers cited above.

might ground reasons to sometimes be coherent, but this is a far cry from always grounding reasons to be coherent. And the intrinsic goodness strategy holds that the intrinsic goodness of coherence grounds a reason to be coherent.

While the intrinsic goodness strategy is the most plausible of the three when it comes to weak significance, it is still somewhat implausible that coherence is intrinsically good. What is *intrinsically* good about it? Moreover, even if coherence did turn out to be intrinsically good, it's not plausible that the reason grounded by such goodness would be very strong, very often. It's cold comfort for those of us who think rationality is very important to find out that we have weak reasons to be rational because coherence is an intrinsic good! These three strategies seem the most prima facie plausible. Since none of them seem promising, it's plausible to think that the wide-scope requirements are neither weakly nor strongly deontically significant.

A recap is in order. I started off by clarifying two ways in which rationality could be deontically significant. It is *strongly significant* if it is such that we always ought to do what rationality requires. It is *weakly significant* if it is such that we always have reason to do what rationality requires. The next step was to investigate what rationality requires. Those in the literature start with the assumption that there is a tight connection between rationality and coherence. There are two main views about what the connection is. The narrow-scope view was rejected because if it were true then it's very plausible that rationality is not deontically significant in either the strong or the weak ways. This left the wide-scope view. Although the wide-scope view is traditionally motivated by the thought that it alone can account for the deontic significance of rationality, it turns out to be mysterious how the wide-scope view can account for the deontic significance of rationality. What we're left with is scepticism about the deontic significance of rationality.

3. A better way: Reasons Responsiveness

I started the discussion of rationality in the previous section by assuming that there is a tight connection between rationality and coherence. This led to the narrow- and wide-scope views, which led to scepticism about the deontic significance of rationality. That is not a good place to be. I think the first thing to do in response to these problems is explore the possibility that rationality doesn't have the type of connection to coherence that narrow- and wide-scopers assume.

Instead, I think we should hold Reasons Responsiveness.

> ***Reasons Responsiveness***: Rationality consists in correctly responding to the objective normative reasons one possesses.[21]

This section has two goals. In the first subsection I will further explicate and motivate Reasons Responsiveness, and defend it against some initial objections. The goal will not be to provide a knock-down argument for Reasons Responsiveness. Instead, the goal will be to show that the view can be motivated and that it can withstand the blow of some initial objections. In the second subsection I will situate the view in the dialectic about the deontic significance of rationality. This will set up this paper's main argument for Reasons Responsiveness, which aims to show that Reasons Responsiveness (unlike its rivals) can elegantly explain the strong deontic significance of rationality.

3.1 Reasons Responsiveness, what and why

Reasons Responsiveness contains two important pieces of ideology that need further elaboration. The first is *objective normative reasons* (herein I will for the most part drop 'normative'). The objective reasons are the reasons constituted by facts. They are, to use the usual phrase of initiation, the facts that count in favour of various reactions—actions, beliefs, desires, etc. The fact that the building is tall is an objective reason not to jump off it, the fact that Anne is at home is an objective reason not to believe she is at the store, the fact that the cheeseburger is delicious is an objective reason to desire it, and the fact that the cheeseburger is high in fat is an objective reason to intend not to eat it.

Not *all* of the objective reasons matter to rationality. This brings me to the second important piece of ideology, namely, the notion of possessing normative reasons. The only reasons that matter when it comes to Reasons Responsiveness are the reasons that you *have* or *possess*.[22] You have some type of epistemic access to the reasons that you possess. It's plausible that only the reasons you possess affect the

[21] Broome (2007a, 2013) considers whether rationality consists in correctly responding to reasons. He argues that it doesn't. As I show in Lord (2014a) and Lord (MS, ch, 2), his arguments don't have any traction against certain views that appeal to the reasons you possess.

[22] Often philosophers will use the phrase '*A* has a reason to ϕ' when they merely mean 'there is a reason for *A* to ϕ'—when they don't mean to imply that *A* stands in any epistemic relation to the reason. From here on in I will only use the possession talk when I mean that one possesses a reason in the sense central to the Reason Responsiveness view. When I merely want to say that there is a reason to ϕ, I will use the 'there is' construction.

rationality of your attitudes and actions. This is plausible, because when some fact is completely outside your ken, it doesn't seem to bear on the rationality of your actions and attitudes.[23]

With the two main pieces of ideology on the table, we can now ask what it means to say that rationality *consists in* correctly responding to reasons. I'll make two points about this. First, to say that rationality consists in correctly responding to reasons is to say that *what it is* to be rational is to correctly respond to reasons. This is intended to be a real definition of the property of being rational. Actions can have this property, beliefs can have this property, desires can have this property, people can have this property (among many other things). Reasons Responsiveness holds that the essence of an action's (or belief's or desire's or person's) being rational is its being a correct response to the possessed reasons.

The second point I will make is that I think that rational requirements are generated by possessed reasons. When a set of possessed reasons *S decisively* supports ϕ-ing, ϕ-ing is rationally required. When a set of possessed reasons *S sufficiently* supports ϕ-ing, ϕ-ing is rationally permitted.[24] This means that Reasons Responsiveness posits a certain class of narrow-scope requirements. These requirements maintain that *if* a set of reasons *S* that *A* possesses decisively supports ϕ-ing, *then A* is rationally required to ϕ.

Now that we have a better understanding of Reasons Responsiveness, let me provide some initial motivation for the view. The view is a version of the natural idea that what it is rational to do is a function of one's information. Despite the intuitive appeal of this idea, views of this type have fallen out of favour in metaethics.[25] This is

[23] I won't defend here any particular view about what it takes to possess a reason. I think that all of the plausible views will be able to explain what's going on in the cases that will interest us. The three most plausible views, I think, are the views that in order to possess some reason *r*, you have to know *r*, justifiably believe *r*, or be in a position to know *r*. In Lord (MS, chs. 3 and 4) I defend the last option.

[24] For more on decisiveness and sufficiency, see Schroeder (2015b) and Lord and Maguire (2016).

[25] There are many views of rationality across philosophy that embrace the basic thought that rationality is a function of one's information. This includes some views in metaethics—most prominently, the views of Parfit and Schroeder. My view is importantly different from these views, because it holds that rationality is a function of the *objective reasons*. Parfit and Schroeder hold that rationality is a function of apparent, or subjective, reasons. These reasons are just a function of one's perspective, *even when it is seriously deluded*. This causes many problems for their views. In this context, it makes it much less plausible that they can vindicate the deontic significance of rationality. They will have bootstrapping problems. To some extent,

because the metaethical literature has turned its attention to coherence requirements. Given this dialectical focus, I will provide a preliminary argument for Reasons Responsiveness by arguing that it has serious advantages over the wide-scope and narrow-scope views.[26]

I'll start by arguing that Reasons Responsiveness can explain the datum that motivates the wide-scope and narrow-scope views. The datum, recall, is that one is irrational if one is incoherent. I think that Reasons Responsiveness can explain this datum. This is because it's plausible that whenever one is incoherent, one fails to correctly respond to some of the possessed reasons.

To use our earlier example, it's plausible that there will never be a single time where I possess sufficient reasons to believe that I am writing and sufficient reasons to disbelieve that I am writing. I will always possess decisive reasons to not have at least one of those beliefs. In our example, it's disbelieving that I am writing. I possess decisive reasons to not have that belief. Thus, when I hold it, I am irrational. Note that, according to Reasons Responsiveness, I am irrational not because I violate one of the narrow- or wide-scope coherence requirements. Instead, I am irrational because I violate a requirement generated by possessed reasons. In the example, I possess decisive reasons not to disbelieve I am writing. Thus, I am rationally required not to disbelieve. I am irrational because I am violating *that* requirement.

My general strategy for showing that whenever you are irrationally incoherent you are failing to correctly respond to reasons—which I have pursued at length in Lord (2014a) and Lord (MS, ch. 2)—is to show that the kind of explanation just given generalizes. While it is very plausible that one cannot have sufficient reason to believe p and sufficient reason to believe $\neg p$ at the same time, it's natural to worry that there are other irrationally incoherent sets of attitudes that one can get oneself into without failing to correctly respond to any possessed reasons. Although this is not the place to provide the argument that the strategy is generally successful, let's explore another paradigmatic case: means–end incoherence.

As we've already seen, it's very plausible that you are irrational whenever you intend to ϕ, believe that ψ-ing is necessary for ϕ-ing, and yet fail to intend to ψ. According to my view, this is true because

Schroeder embraces this by embracing a narrow-scope account of rationality. I argue elsewhere (Lord MS, ch. 3) that this leads to serious problems for their accounts of possession as well.

[26] This is a condensed version of the argument I provide in Lord (2014a) and Lord (MS, ch, 2).

whenever you are means–end incoherent, you fail to correctly respond to the reasons you possess. To see why this is plausible, let's return to John the Stilton-lover. John intends to eat a planet made of Stilton and believes that in order to do this he needs to intend to build a spaceship, yet he fails to intend to build a spaceship. Reasons Responsiveness rightly predicts that John is irrational. Why? Most likely it's because he possesses decisive reasons *not* to intend to eat a planet made of Stilton. So when he does intend to do that, he is irrational.

Of course, it won't always be the case that one's end is irrational. So, to take another case, suppose that I intend to revise this paper, believe that in order to do so I must intend to turn on my computer, yet fail to intend to turn on my computer. My end in this case is rational (and we can assume my belief is too). So why am I irrational in this case? The reason I am irrational is that I have decisive reason to intend to turn on my computer and I don't. I have decisive reason to do this because the weight of reasons to intend some end transmit to reasons to intend the necessary means to those ends.[27] Given this transmission principle, possessed decisive reasons to ϕ will guarantee that there are possessed decisive reasons to take the necessary means to ϕ-ing.[28]

As has been pointed out in the literature, there is a further kind of case that the above story doesn't handle neatly.[29] This is a case where one has sufficient but not decisive reasons to intend some end ϕ. Consider my reasons to revise this paper right now. I have sufficient reasons to do this, but they aren't decisive—I could permissibly do something else instead. Suppose I do intend to revise right now, and I know that in order to revise I have to type but I do not intend to type. I am means–end incoherent. But I cannot use the story just told to explain why I am violating a requirement generated by possessed reasons. This is because the story I just told relied on a transmission principle about reasons that are *decisive*. The reasons to intend the end aren't decisive in this case. So I need a different story to explain this case.

[27] For discussion and defence of the basic idea, see Kiesewetter (2015).

[28] Although I'll skip it to save space, there is obviously a third way one can go wrong: one can have an irrational means–end belief. In some cases of means–end incoherence, the rational failure will be to have an irrational means–end belief.

[29] See Kolodny (2007, 2008a), Way (2012, 2013), and Lord (2014a; MS, ch. 2).

I think that what is going on in this case is that when I form the intention to revise, the weight of my reasons to intend to type are *intensified* to the point of being decisive. Thus, once I form the intention to revise, I come to have decisive reason to intend to type. In many cases my reasons to intend to revise remain merely sufficient. This explains why it often remains permissible for me to revise my end. Still, when I intend the end, my reasons to take the means are decisive. As I argue in Lord (2014a) and Lord (MS, ch, 2), this story allows us to explain data that have proven very hard to explain, namely, (i) why it is that something is always going wrong when one is means–end incoherent, (ii) why one is always required to take the means when one intends the end in these cases, and (iii) why one is often permitted to drop one's end even though one is required to take the means (given that one intends the end).[30]

One might doubt my story about coherence by appealing to cases where it looks as though one has decisive reasons to be incoherent. For example, an eccentric billionaire might offer me $1 million to intend to revise the paper, believe that in order to revise the paper I must intend to turn on the computer, and fail to intend to turn on the computer. Wouldn't this offer provide me with decisive reason to be incoherent? Isn't that incompatible with my failing to correctly respond to the possessed reasons whenever I'm incoherent?

In order to deal with this problem, one needs to draw the distinction between the *right kind* and *wrong kind* of normative reasons for attitudes. The right kind of reasons for attitudes are the reasons that explain why some attitudes are *correct* or *fitting*. The wrong kind of reasons merely incentivize having attitudes, whether they are fitting or

[30] Objection: Imagine a case where the only attitude one has is the belief that one ought to ϕ. In this case, one is incoherent by being akratic. However, it's hard to see how one could possess decisive objective reasons to either give up the belief or form the intention to ϕ. Thus, it's hard to see how my strategy can be extended to this case. Response: This case definitely causes problems for my strategy if we assume that in order to possess reasons, we need to have attitudes that have those reasons as their contents (call these Holding Views). Even if a Holding View is true, it's not clear how problematic such an unusual case is. I am not particularly worried if the strategy breaks down at the extreme limit. That said, it's not clear that this case ruins the strategy if we reject Holding Views. As it happens, at the end of the day I do reject Holding Views (cf. note 23) in favour of the view that possession requires being in a position to know. If possession only requires being in a position to know, it's not clear if this case ruins the strategy, for one could be in a position to know decisive reasons to drop the belief or decisive reasons to form the intention, even if one only has one attitude. My hunch is that this is how all filled-in versions of the case would go. Thanks to an anonymous referee for bringing this case to my attention.

not.[31] To get a handle on the distinction, consider Pascal's wager. The fact that one will experience infinite rewards in heaven only if one believes that God exists provides the wrong kind of reason to believe that God exists. The only things that provide the right kind of reasons for believing that God exists are pieces of evidence for God's existence. This is because only evidence can explain why a belief is fitting or correct. The wrong kind of reasons can only explain why it's fortunate that one has a particular belief.

The current problem can be solved by amending Reasons Responsiveness so that only the possessed *right kind* of reasons are relevant to rationality. The billionaire's promise provides the wrong kind of reason to be incoherent. The right kind of reasons for intention are a function of the features of the *intended acts* (rather than a function of the potential benefits of having the intention). Given this, it's plausible that the right kind of reasons for intention will either decisively support intending to revise, and thus also support intending to turn on the computer, or decisively support not intending to revise. Either way, if I am means–end incoherent, I am failing to respond to some of the possessed right kind of reasons.

This move will be ad hoc unless there is independent motivation for restricting the class of reasons to reasons of the right kind. Fortunately, there is independent motivation. We can see this by reflecting on the usual reaction to Pascal's wager. Most take Pascal's wager to be a very bad defence of belief in God. This is because it doesn't appeal to the evidence. We can provide a very nice explanation for why Pascal's wager is a bad defence of belief in God by appealing to the claims that only the right kind of reasons are relevant to rationality and that incentives to have beliefs are reasons of the wrong kind. This provides independent motivation for thinking that only reasons of the right kind are relevant to epistemic rationality.[32]

Assuming the preceding story can be generalized, Reasons Responsiveness can explain the data that motivate the narrow- and wide-scope views. It also enjoys significant advantages over each. Start with the narrow-scope view. One feature of the narrow-scope view is that it requires people to have particular attitudes. This is a feature

[31] The last two sentences are not meant to be analyses of the right and wrong kinds of reasons. It is, as one should expect, very controversial what the correct analyses are. For my preferred way of thinking about things, see Schroeder (2010) and Lord and Sylvan (MS).

[32] As the literature on the right kind of reasons has made abundantly clear, an analogous story can be told about reasons for intention by appealing to Kavka (1983)'s toxin puzzle.

because it is plausible that we are sometimes rationally required to have particular attitudes. If I rationally intend to ϕ and I know that in order to ϕ I have to intend to ψ, it's plausible that I'm rationally required to intend to ψ. Unfortunately, the way in which the narrow-scope view generates this feature is buggy. This is because it leads to bootstrapping. The narrow-scope view makes no distinction between rational ends and irrational ones. This leads to the prediction that even when I have completely delusional ends and means–end beliefs, I'm required to intend the means.

The wide-scope view removes this bug by giving up on the feature altogether. It removes the bug by insisting that no particular attitudes are ever required.[33] All that is required is that one avoid certain incoherent combinations. Reasons Responsiveness, it seems, can remove the bug without giving up on the feature. It doesn't lead to bootstrapping, because the cases that cause trouble for the narrow-scope view are all cases where one's antecedent attitudes are unreasonable. In those cases, Reasons Responsiveness will not predict that one is required to intend the means. It will predict that one is required to give up the antecedent attitudes. On the other hand, Reasons Responsiveness does predict that one is often required to have particular attitudes. It predicts that when one rationally intends to ϕ and knows that in order to ϕ one has to intend to ψ, one is required to intend to ψ. This seems like the right result.

Thus, Reasons Responsiveness can be motivated by considering the dialectic between the narrow- and wide-scoper. It seems as though it can explain the data that motivate those views, while avoiding some of their bigger flaws. The case for the view will be even stronger once we combine this result with the main result of this paper—that Reasons Responsiveness is the only view that vindicates the strong deontic significance of rationality.

Before turning to how Reasons Responsiveness bears on the deontic significance of rationality, let's consider two further objections. The first has to do with cases involving rational but false beliefs. Take Williams's famous gin and tonic case (1980). Suppose Bernie is at a normal party at a bar. He orders a gin and tonic from the bartender. The bartender gives him a glass of clear liquid. He believes—as any sane person in his situation would—that the glass contains gin and tonic. He forms the intention to drink the liquid. This intention seems

[33] This has led many to object to the so-called symmetry of the wide-scope requirements. See Kolodny (2005), Schroeder (2004, 2009b), Bedke (2009) and Lord (2014b).

rational. However, unbeknownst to him, the glass really contains petrol.

Objective normative reasons are constituted by facts. It is not a fact that the glass contains gin and tonic. So it can't be, on my view, that the reason that rationalizes Bernie's intention to drink is the fact that the glass contains gin and tonic. So it looks as though my view cannot explain why Bernie's intention is rational.

It is true that Bernie's intention is not rationalized by the fact that the glass contains gin and tonic. However, that doesn't mean Bernie doesn't possess objective reasons to intend to drink. To see this, think of what rationalizes his belief that there is gin and tonic in the glass. It seems as though there is a host of facts that rationalize his belief—for instance, the fact that he's at normal bar at a normal party and just ordered a gin and tonic from the bartender. He possesses these reasons, and they plausibly play a role in rationalizing his belief. These reasons are also, I claim, objective normative reasons he possesses to intend to drink. Moreover, since they sufficiently support the claim that the glass contains gin and tonic, they also sufficiently support forming the intention. Thus, my view can explain why Bernie's intention is rational.[34]

I should be explicit that making this move requires one to think that reasons come cheaper than most ethicists think they do. The main reason why ethicists think that cases like Bernie's provide such a challenge for a view like mine is that they are disposed to think that the only thing that could possibly be the reason that rationalizes Bernie's intention is that the glass contains gin and tonic.[35] They have a hard time thinking that the reasons that rationalize Bernie's belief are objective reasons.

Before I explain why I think they are objective reasons, let me introduce the second—very much related—objection. Another batch of examples that some think cause problems for Reasons Responsiveness have to do with testimony. To take the hardest case, imagine that a trusted and reliable adviser tells you that you shouldn't go into the next room, without telling you why. It seems rational for

[34] I first defend this explanation in Lord (2010). See also Lord (forthcoming a) and Lord (MS, ch. 7).

[35] Some good evidence for this is the way the discussion goes in Schroeder (2008). Although he considers a few other views, the one that he takes to be the most plausible is the view that Bernie's reason is that the glass contains gin and tonic.

you not to go into the next room. According to some, my view has trouble explaining why. This is because many think that the fact that you shouldn't go into the next room cannot be an objective reason not to go into the next room, nor can the fact that someone told you that you shouldn't go into the next room.[36] But these facts seem to be the only possible contenders for being the reasons you possess not to go into the next room. If they aren't objective reasons, then it seems as if my view cannot explain why you are rational when you avoid the room.

Needless to say, I think that both of those facts—the existential fact about what you should do and the testimonial fact—can be objective normative reasons. And this is for the same reason that I think the reasons that rationalize Bernie's belief can be objective normative reasons to intend to drink. The reason why I think these facts are objective normative reasons is that those facts bear the earmarks of objective normative reasons.[37]

First, they intuitively count in favour of the actions and attitudes. You should treat what the adviser says as a reason in deliberation. Indeed, you should treat it as a very strong reason. Similarly, you should treat the reasons you have for thinking the glass contains gin and tonic as reasons to drink what's in the glass. And if those reasons are strong enough to make it rational for you to believe the glass contains gin and tonic, you should treat those reasons as weighty reasons to intend to drink. Because of this, it's plausible that these facts can justify actions and attitudes. They can be the things that explain why it makes sense to pursue a particular course of action.

Moreover, they can also be the reasons *for which* we act or hold attitudes. You can, and plausibly do, refrain from going into the next room because you shouldn't. Moreover, you can refrain from going in because someone told you that you shouldn't.[38] The reasons you have

[36] See, for example, Broome (2008b), McKeever and Ridge (2012), and McNaughton and Rawling (2011).

[37] Furthermore, I think that stock objections to the idea that they can be objective reasons are misguided. Some objections involve misguided worries about double- counting (see Schroeder 2009a for reasons why these objections are misguided). Others appeal to intuitions about right-making, holding that these facts can't intuitively be right-makers (see, for example, Broome 2008b, McKeever and Ridge 2012, and McNaughton and Rawling 2011). This thought cuts both ways, however, for it does seem as though these facts can justify actions and attitudes. Given this, it's unclear why we shouldn't think that these reasons are counterexamples to whatever theory of right-making objectors have in mind.

[38] Compare this with the arguments given in Schroeder (2009b) for the claim that existential facts about what you ought to do are themselves reasons.

to believe the glass contains gin and tonic can also be (among) the reasons for which you intend to drink. Indeed, it's quite plausible that they are in fact among the reasons you intend to drink. Evidence for this comes from the fact that those would be the natural facts for you to turn to in order to explain yourself if someone pointed out that the glass didn't contain gin and tonic.

Finally, these facts can be defeated. If another reliable adviser tells you that you should go into the next room, this defeats the reason provided by the other adviser's testimony. It seems as though, in this case, there isn't sufficient reason to go in or not to go in (one should seek more information). Moreover, it wouldn't be rational for you to decide about whether to go in or not (again, the rational thing to decide to do is seek more information). Similarly for the reasons to believe the glass contains gin and tonic. If you get evidence that the glass doesn't contain gin and tonic or you receive information that makes you doubt the reliability of your reasons, you cease to have good enough reason to intend to drink.

Counting in favour, being the things we can act for, and defeasiblity are core features of objective normative reasons. I think that it is extremely plausible that a lot more facts have these features than many ethicists let on. It's true that some of the theories that some of these ethicists hold predict that reasons don't come very cheaply. So much the worse for their theories. At the very least, they have to explain why the types of facts we're discussing here aren't objective reasons.

I don't pretend that this is a full defence of Reasons Responsiveness. It is only a preliminary defence. I have provided a much fuller defence of the points made in this section elsewhere (Lord 2010, 2014a, forthcoming a, MS). I remind the reader now that *the main argument in this paper* for Reasons Responsiveness is anchored in the claim that it can vindicate the deontic significance of rationality. With this firmly in mind, let's return to the deontic significance of rationality.

3.2 Reasons Responsiveness and the deontic significance of rationality

Before moving on to how Reasons Responsiveness handles the deontic significance of rationality, it is worth pausing to attend to a subtlety of Kolodny's views.[39] One might think that Kolodny is not exploring the deontic significance of rationality per se, but rather just the deontic significance of coherence requirements like *Means–End N* and *Means–End W*. He lends this impression at the very beginning of Kolodny

[39] I thank an anonymous referee for pressing me to say something about this.

(2005) when he discusses two different ways of talking about rationality. According to one way of talking, what rationality requires is determined by normative reasons—he calls this objective rationality. According to a different way of talking, what rationality requires is determined by considerations about coherence—he calls this subjective rationality. He then makes it clear that he is only interested in discussing subjective rationality. It is plausible to conclude from this that the only way to refute Kolodny's arguments is by showing that the coherence requirements are deontically significant. This is not the tack I will take. As we saw in §2.2, I agree with Kolodny that the coherence requirements are not deontically significant. The tack I will take is to argue that the requirements that are determined by the reasons one possesses are deontically significant. Thus, if this interpretation of Kolodny is right, my arguments appear not to be responsive to his.

I will make three points about this. The first is that Kolodny's discussion is simplistic in an important way. As he is thinking of it, the contrast is between a fully objective notion of rationality—one that is tied to all of the facts—and a fully subjective notion of rationality—one that is tied solely to one's non-factive internal states. This contrast neglects my kind of view, which places importance both on the mind-independent reasons and on the agent's perspective. Thus, I don't think we can conclude that Kolodny is not interested in my notion of rationality simply because he denies interest in his objective notion of rationality.

Still, this first point doesn't show that Kolodny and I are talking about the same concept. My second point is that I think Kolodny and I are using the same concept, because we both agree on some basic roles that the relevant concept of rationality is supposed to play. While this might not be a deductive proof that we are using the same concept, it is good evidence. What roles do we agree the concept plays?

First (and perhaps foremost), we agree that the relevant concept of rationality provides a distinctive kind of explanation of why there is something amiss with (at least some forms of) incoherence.[40] Kolodny maintains that certain narrow-scope coherence requirements are the materials that the correct theory of this concept uses to explain what is amiss with incoherence.[41] I deny this. Nevertheless, as we saw

[40] See, for example, Kolodny (2007, §4.2).

[41] The content of the requirements Kolodny endorses changes slightly between papers, but the basic idea is that you are rationally required to ϕ if you believe you have decisive reason to ϕ.

in the last subsection, I think that we can explain why there is something amiss by appealing to my theory's requirements.

A second role for the concept is its connection to a certain sort of criticism.[42] When someone is irrational, they are open to a particular kind of criticism. This is because there is some part of their perspective that severely clashes with some reaction of theirs. This doesn't necessarily mean that *every* time someone is irrational they should be condemned; it just means that a certain kind of criticism is apt in paradigm cases of irrationality. This is a role that I think my concept of rationality plays. I think the concept Kolodny is picking out is also supposed to play this role.[43]

As I said above, while this is good evidence we are using the same concept, it doesn't provide a deductive proof. My third point is that even if Kolodny is stipulating that the concept of rationality he is interested in is necessarily tied to the coherence requirements, my argument is important for the debate he is engaging in. This is because if he is so stipulating, the debate should be seen as a debate about what David Plunkett, Alexis Burgess and Timothy Sundell call *conceptual ethics*.[44] In other words, the debate is not about the nature of some concept, but rather about which concepts we should use. If it is a debate about this, and my arguments below are sound, then I think my concept has a huge advantage over Kolodny's. This is because my concept is intimately connected to the concept of deliberative obligation. Given this connection, my concept can vindicate many more of the platitudes about rationality than Kolodny's concept can—namely, all of the platitudes associated with deontic significance. Thus, even if Kolodny and I are using different concepts, there is a debate to be had between my view and Kolodny's. Further, there are strong reasons to think my view is winning that debate.

[42] A further consideration in favour of the hypothesis that Kolodny and I are talking about the same concept is that we are both hypothesizing about Broome's concept, which is centrally tied to the requirements of rationality being *strict* (see especially Broome (1999). The strictness of the requirements explains why one is open to criticism, and explain why there is something distinctively amiss when one is incoherent. Further, Broome clearly thinks that *it's possible* that his concept is analysed in terms of reasons. After all, he has written several papers about whether rationality can be understood in terms of reasons (see Broome 2007a, 2013, and Lord 2014a for an antidote). Although he argues against the view, it's not on the grounds that it invokes a separate concept.

[43] It is uncontroversial that Kolodny's notion of objective rationality does *not* play this role. So this is a further way of distinguishing my notion from that one.

[44] See Burgess and Plunkett (2013a, 2013b), Plunkett and Sundell (2013), and Plunkett (2015).

Now let's turn to the deontic significance of Reasons Responsiveness. Note that it follows immediately from Reasons Responsiveness that rationality is weakly deontically significant. This is because in order to be rationally required to ϕ, I must possess objective reasons to ϕ. Thus, there will always be objective reasons to ϕ when rationality requires me to ϕ. Securing the weak deontic significance of rationality is thus very easy for my view.

Nevertheless, it's far from clear that this view can vindicate the strong deontic significance of rationality. Even though the view entails that there will always be reasons to do what rationality requires, it's far from clear that these reasons will always be weighty enough to ground an obligation. Moreover, according to a very popular view about what one ought to do, it's obvious that sometimes the reasons one possesses won't be sufficiently strong to make it the case that one ought to do what's rationally required.

According to this view, which we'll call *objectivism*, what you ought to do is determined by all the objective reasons. You ought to ϕ, according to the objectivist, only if the balance of all the reasons decisively supports ϕ-ing. If objectivism is true, then Reasons Responsiveness cannot vindicate the strong deontic significance of rationality. This is clear by reflecting on cases. Here's the case from the introduction:

> *Sick Mother*
> Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling art historian/dealer, doesn't have the money. It looks as if his mother is going to die. She will, however, be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go to see her. As it happens, a pawnshop owner is Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. The rub, of course, is that he has no idea that this pawnshop even exists, much less that such a deal is to be had there.

In *Sick Mother*, Jack possesses good reasons to go to California—his mother will be greatly comforted by his presence. However, there is also a very good reason for him to go to Queens—the fact there is a cheap Picasso to be had. The objectivist says that, in this case, Jack ought to go to Queens, even though it seems that it would be

irrational for him to go to Queens. It would be irrational because the reasons he possesses strongly support his going to California.

This is a serious challenge. In order to see if it can be met, we must look at what the best theory is of what we ought to do. If objectivism is the best theory, then Reasons Responsiveness cannot vindicate the strong deontic significance of rationality. This would be a big blow to those of us who think rationality is deontically significant. In fact, I don't much care about a mere vindication of the weak significance of rationality. I will only be fully satisfied if rationality turns out to be strongly deontically significant.

The rest of this paper will aim to meet the challenge. I don't think that objectivism is the best theory of what we ought to do. Instead, I think that the best theory holds that what we ought to do is what is decisively supported by the reasons we possess. I think this view can be well motivated independently of any debate about rationality. This, I think, is as it should be if rationality really is strongly significant.

4. Ignorance and obligation

In this section I will argue for *Possessed Reasons*:

> ***Possessed Reasons*:** What one ought to do is determined by the reasons one possesses.[45]

[45] One of the major objections to anti-objectivist views appeals to bystanders with more information. It seems as if those with more information can have true thoughts about what one ought to do that isn't what one ought to do relative to what one knows. My goals in this paper do not include fully answering this objection. My main goal is to provide a new positive argument for *Possessed Reasons*. However, I think that this problem can be solved. My solution has two parts. The first part is to show that *Possessed Reasons* is compatible with the thought that deliberation aims at what's best (or what's supported by all the reasons). *Possessed Reasons* is compatiable with this claim because correct pursuit of one's aims might be constrained in certain ways. Here are two examples. Our obligations might be constrained by our physiological abilities even though we aim to do what's best. Our epistemic obligations might be constrained by, for example, the evidence, even though the aim of epistemic deliberation is to believe the truth. It doesn't follow from the claim that the aim of epistemic deliberation is to believe the truth that we always ought to believe the truth. The second part of the solution appeals to the semantics of 'ought' in English. 'Ought' is a flexible word in so far as we can relativize our 'ought' claims to different bodies of information. This means that advisers with more information can have 'ought' thoughts about our obligations that are relativized to their information. This is what they're doing, I claim. It's right that this means they aren't having thoughts about our deliberative obligations per se. This is unsurprising, though, given that the aim is to do what's best. So even though they aren't thinking about our deliberative obligations, they are having thoughts that are relevant to the deliberative project. See Lord (2015) for a defence of this response.

I will first discuss some potential counterexamples to objectivism. With these cases in mind, I will turn to my main argument for *Possessed Reasons*. This argument runs through the claim that if the members of some set of reasons *S* make it the case that you ought to ϕ, then you can ϕ because of the members of *S*. I'll argue that you can only ϕ because of the members of *S* if you possess them. If sound, this argument will establish that a necessary condition for some set of reasons *S* making it the case that you ought to ϕ is that you possess the members of *S*. I will end this section by arguing that once you accept this necessity claim, the relevant sufficiency claim will follow from very plausible assumptions.

4.1 Neighbours and envelopes

The purpose of this section is to discuss two kinds of case that have been thought to provide counterexamples to objectivism. The dialectic about these cases will provide helpful set-up for my main argument for *Possessed Reasons*.

It is obviously true that we are almost always ignorant of the full effects our actions will have. Take this example, analogous to *Sick Mother*, discussed by Thomson (1990) and Scanlon (2001, 2008).[46]

> *Day's End*
> Jack always comes home at 9:00 p.m., and the first thing he does is flip the light switch in his hallway. He did so this evening. His flipping the switch caused a circuit to close. By virtue of an extraordinary series of coincidences, unpredictable in advance by anybody, the circuit's being closed caused a release of electricity (a small lightning flash) in his neighbour's house next door. Unluckily, his neighbour was in its path and was therefore badly burned.

Many, and I am one of them, have thought that cases like *Day's End* strongly support the rejection of objectivism. After all, the objectivist thinks that Jack ought not to flip the switch in his hallway. This is because the fact that it will lead to his neighbour being badly burned is an objective reason not to. Relative to all the reasons there are, this reason is decisive. But since flipping the switch will lead to his neighbour being burned only because of 'an extraordinary series of coincidences' that he has no way of knowing about, it seems quite plausible that, to use Scanlon's words, 'If it is true that [Jack] ought

[46] To be clear, I do not think that *Day's End* is different in kind from *Sick Mother*. I introduce *Day's End* because I would like to quote from Scanlon's discussion of it.

not to have flipped the switch, this is true only in a sense of "ought not" that seems to me to lack the moral content that the idea of permissibility has. Both [Jack] and [his neighbour] may wish, after the fact, that [Jack] had not flipped the switch, but in doing so [Jack] did not act impermissibly' (Scanlon 2008, p. 48).

Unfortunately for foes of objectivism, there is a standard objectivist reply to this line of reasoning.[47] The core insight is that we must cleave apart the deontic facts—facts about what ought to be done—from the hypological facts—facts about blame and praise. Once we do this, the response goes, we can see that arguments like Scanlon's against objectivism try to draw conclusions about the deontic from conclusions about the hypological. They get their bite only if we assume that the fact that Jack is not blameworthy for flipping the switch entails that he acted permissibly. Once we give up the idea that permissibility lines up with blamelessness, the argument falls apart. Jack might be blameless, but he still does what he ought not to do. Indeed, it seems that there is a good explanation of why Jack is blameless despite doing wrong in this case. Jack is blameless because it is rational for him to believe that flipping the switch is permitted by all the reasons. But his belief is *false*. This is what explains why he is still doing something wrong.

I agree that we must separate the deontic facts from the hypological facts.[48] Moreover, I agree that cases like *Day's End* lose their initial dialectical bite once this standard move is made. Luckily, though, there are other well-known cases where the standard move is less than effective. *Three Envelopes* is a case with this structure:[49]

> *Three Envelopes*
> Suppose Margaret is given the choice of picking one of three envelopes placed in front of her. Margaret is informed that the third envelope contains $900. She is also informed that either the first envelope or the second envelope contains $1000, and that whichever envelope doesn't have the $1000 in it is empty. So, given her evidence, there is a 0.5 chance that the first envelope contains $1000 and a 0.5 chance that the second envelope contains $1000.

[47] This move goes back to Moore (1912). Thomson herself replies this way in Thomson (1990, 2008). See also Graham (2010).

[48] Many other anti-objectivists do this as well. See especially Scanlon (2008) and Zimmerman (2008).

[49] This particular case initially comes from Ross (2006). It is discussed further in Schroeder (2009b) and Ross (2012).

Intuitively, Margaret ought to choose the third envelope—the one she knows has $900 in it. Moreover, there is an important difference between *Three Envelopes* and *Day's End*, namely, that in *Three Envelopes* Margaret *knows* that choosing envelope three is not the best option. That is, she knows that it is not the option decisively supported by all the reasons. The option best supported by all the reasons is choosing the envelope with $1000 in it. Despite the fact that Margaret knows that choosing envelope three is the second-best option, it still seems as though she ought to choose it.

This difference between *Day's End* and *Three Envelopes* makes the standard objectivist move much less plausible when it comes to *Three Envelopes*. For in *Day's End*, it's rational for you to think that flipping the switch is permitted by the balance of all the reasons. It's not rational to think that choosing envelope three is permitted by the balance of all the reasons in *Three Envelopes*. In fact, Margaret *knows* that it's not. Still, it seems as though Margaret ought to choose envelope three. This is important, because the standard move seems plausible in *Day's End* only because you have a rational yet false belief about what the balance of all the reasons supports. The rationality of this belief helps explain why you are blameless, and the falsity explains why what you did was actually impermissible. *Three Envelopes* doesn't have this structure. Margaret seems blameless all right, but she lacks the false belief. The fact that it still seems that she should choose the second-best option despite the lack of the false belief strongly suggests that she is blameless, because choosing the third envelope is what she ought to do.

Before moving on, let me mention a common reaction to *Three Envelopes* in order to set it aside. Many ethicists' first reaction is to draw a distinction between what you objectively ought to do and what you subjectively ought to do.[50] They then use this distinction to explain intuitions. First, there is the intuition that what Margaret ought to do is choose the envelope with $1000 in it, because that's what she objectively ought to do. Second, there is the intuition that what Margaret ought to do is choose envelope three, because that's what she subjectively ought to do. This is supposed to resolve the puzzle.

[50] This move is made in response to this case in at least Schroeder (2009b), Smith (2008) and Ross (2006).

The problem is that it doesn't resolve the puzzle.[51] At best it changes the subject. That is, at best we find out that there is some sense in which Margaret ought to choose envelope three. It turns out that there are independent reasons to believe in such a sense, and the explanation it gives for why, in some sense, Margaret ought to choose envelope three is a principled one. But this is not the question we were asking. We were asking whether Margaret ought to choose envelope three. We were asking whether 'Yes' would be the correct answer to a question Margaret might ask herself: 'Ought I to choose envelope three?'

The important point is that mere appeal to the distinction between subjective and objective oughts doesn't answer our question, because it doesn't specify which obligation provides an answer to this central deliberative question. If Margaret is conceptually sophisticated enough, she is in a position to know that she subjectively ought to choose envelope three and that she objectively ought to choose envelope one or envelope two. But even if she did know those facts, she could meaningfully wonder which requirement she ought to satisfy. She thus doesn't seem to learn what she ought to do, full stop, by learning what she subjectively and objectively ought to do. There is another question that hasn't been answered yet: what ought Margaret to do? When we theorize about what answers this question, we theorize about the deliberative ought.[52]

The preceding dialectic has convinced many—including me—that objectivism is false. Nevertheless, counterexamples can only do so much. These counterexamples on their own don't tell us why objectivism fails. In the next subsection I will present my main argument against objectivism in favour of *Possessed Reasons*. This argument does provide an explanation for why objectivism fails.

[51] Many—both objectivists and anti-objectivists—have pointed this out. See, for example, Jackson (1991), Graham (2010), Kiesewetter (2011), Kolodny and MacFarlane (2010), Björnsson and Finlay (2010), and Lord (2015).

[52] I hasten to add that there are at least two ways to use the objective–subjective distinction to answer our question. The first way is to insist that what you ought to do is just what you subjectively ought to do. The second way is to insist that what you ought to do is what you objectively ought to do. I obviously don't think that what you ought to do is what you objectively ought to do. But for all I've said, it might be that what you ought to do is what you subjectively ought to do. In fact, it might be that the theory I give about what you ought to do is the correct theory about what you subjectively ought to do.

4.2 Ignorance and acting for the right reasons

I'll start with an argument for the following claim: if the members of some set of reasons *S* make it the case that you ought to ϕ, then you possess the members of *S*.

Let's start with the notion of acting for reasons. When you perform actions for reasons, your acts are explained in a certain way by those reasons. Consider some examples. I put on my coat this morning because it was cold outside—that is, the reason for which I put on my coat was the fact that it was cold outside. I believe the heat is on in my office because it's warm in my office and cold outside. I'll go home tonight at five because dinner is at six. In each of these cases, my action is explained by some fact. Moreover, these explanations are not merely causal (although I think they are causal). They also have a normative element to them. The explanantia are not merely the causes of the actions and beliefs, they are also the justifiers of them.

Now consider *Day's End* again. In *Day's End* you face the choice between flipping on the hall light or not. There are good reasons for you to switch it on and a good reason for you not to switch it on. Moreover, in *Day's End* you can perform the relevant action—you have the ability to refrain from flipping the switch. Not flipping the switch can be an action of yours. So there is a reason for you not to flip it on, and you can perform that action. Thus, the objectivist thinks that you ought to refrain from flipping the switch. The reason why is that it will badly burn your neighbour. That fact is a decisive reason not to flip, according to the objectivist.

There's a catch, though. Although you can refrain from flipping the switch, it *doesn't* seem that you can do so *because* flipping the switch will badly burn your neighbour. It doesn't seem that that could be your reason for acting in the right kind of way. For one thing, you don't believe that flipping the switch will harm your neighbour. You've never even considered that, and rationally so. Given this, there is no way that you could refrain from flipping the switch because flipping it will harm your neighbour.

I think something even stronger is true. I think that even if you did irrationally believe that flipping the switch would badly burn your neighbour and decided to bumble around in the dark, you wouldn't have refrained because flipping the switch would badly burn your neighbour. You'd be acting like an idiot if you did that. We might be able to give some psychological explanation of what you did that cited your belief that flipping the switch would badly burn your

neighbour, but we couldn't say you did it because flipping the switch would badly burn your neighbour. This is because you have no idea that flipping the switch will burn your neighbour.

To make this more plausible, compare Delusional Andy to Surprised Andy:

> *Delusional Andy*
> Andy knows that his wife has always been an extremely loyal person. He also knows that he has no reason to think that she is cheating on him (nor does he have any other reason to end the relationship). Despite this knowledge, he does believe that she is cheating on him. He thus moves out and files for divorce. In fact, his wife is cheating on him.
>
> *Surprised Andy*
> Andy knows that his wife has always been an extremely loyal person. However, much to his surprise, he learns that she is cheating on him—her best friend tells him, he finds some love letters, and he catches his wife with her lover. He thus moves out and files for divorce.

In both *Delusional Andy* and *Surprised Andy*, Andy reasons from a belief that his wife is cheating on him to an intention (and subsequent action) to move out and file for divorce. Despite this similarity, it's plausible that only Surprised Andy moves out and files for divorce because his wife is cheating on him. Delusional Andy is, well, delusional. He has no reason to think that she is cheating on him. He gets completely lucky. Because of this, it doesn't seem as though he files for divorce because she's cheating on him.[53]

Notice that you are a lot like Delusional Andy when you more or less arbitrarily believe that flipping the switch will badly burn your neighbour. It turns out that you are right, but you are just lucky. It's not that that reason is *guiding you*. You have about as little contact with that reason as one can have while believing the proposition that constitutes it. (We'll return to Delusional Andy and Surprised Andy in a moment.)

I think the fact that you can't refrain from flipping the switch because it will badly burn your neighbour *explains* why it's not the case that you ought to refrain. In other words, I think that (1) is plausible:

(1) If the reasons in some set S make it the case that you ought to ϕ, then you can ϕ because of the members of S.

[53] For similar arguments for a similar conclusion, see Hyman (2006), Hornsby (2008), Gibbons (2001) and Marcus (2012). The most visible detractor is Jonathan Dancy (see, for instance, Dancy (2000)).

It's important to point out that if you knew that flipping the switch would burn your neighbour, then you could refrain for that reason. It's just that given how you are when you walk in the door, you aren't in a position to refrain for that reason. Thus I think we should understand the relevant sense of 'can' as relativized to one's epistemic position. A stab at the relevant necessary condition: one can ϕ because of r only if one's epistemic position doesn't bar one from ϕ-ing because of r.[54] It's plausible that a sufficient condition for one's epistemic position barring one from ϕ-ing because of r is one's epistemic position barring one from knowing r. This might be necessary too. It all depends on the correct analysis of ϕ-ing because of r. I don't need to give such an analysis, though.[55] All I need here is this sufficient condition.

What's the difference between the cases where you can ϕ because of some reason r and cases where you can't? My answer is that the cases where you can ϕ because of r are always cases where you *possess* r. You don't possess r when you are (non-culpably) ignorant of r, and thus can't ϕ because of r when you are (non-culpably) ignorant of r.[56] When you combine this with (1), you get an argument for the claim that possessing the members of a set of reasons S is a necessary condition for those reasons making it the case that you ought to ϕ.

(1) If the reasons in some set S make it the case that you ought to ϕ, then you can ϕ because of the members of S.

(2) If you can ϕ because of the members of S, then you possess the members of S.

(C) If the reasons in some set S make it the case that you ought to ϕ, then you possess the members of S.

If this argument is sound, then objectivism is false—*Day's End* is a counterexample. So objectivists need to deny one of the premisses. In fact, objectivists need to resist what I said about (1). Objectivists have

[54] There are obviously other necessary conditions—for example, one has to have certain physiological abilities.

[55] See Lord (MS, chs. 5 and 6) for an analysis.

[56] Not all ignorance is created equal. If we hold that A is ignorant about p just in case A fails to believe p, then I think it's implausible that you always fail to possess the reasons you are ignorant of. This is why in Lord (MS, ch. 3) I argue for a *non-holding* account of possessing reasons. Non-holding accounts claim that you can possess reasons you don't believe. However, I think you are culpable for the reasons you possess that you are ignorant of. Hence the qualification in the text.

two choices when faced with (1). They can either deny it or they can argue, *pace* what I argued above, that you can act for the reasons that require you to ϕ when you're ignorant of those reasons. I think both options are implausible. In order to see this, it's important to consider the connection between ϕ-ing because of the reasons that require you to ϕ and your action being creditworthy.

Suppose I ought to buy Anne a hat for her birthday because it would make her happy. And suppose I do buy her a hat. These conditions are not sufficient for my action to be creditworthy. In order for my action to be creditworthy, there must be the right kind of connection between the reasons that justify the action and my performance of the action. To illustrate, my act would be creditworthy if I were to buy her a hat because it would make her happy. On the other hand, if I bought her the hat only because it would make *me* happy (it would cover up her hair, which I find distasteful), it doesn't seem that my act would be creditworthy. It would only be an accident that the following conjunction is true: I bought the hat and the hat makes Anne happy. If it's merely an accident that the act I actually perform is the particular act I ought to perform, then it doesn't seem that I deserve credit for it.

This fact is partly explained, I think, by the fact that when it's an accident that the act one performs is the act one ought to perform, one doesn't act for the right reasons—the reasons that make it the case one ought to do that thing.[57] It's plausible that there is some tight connection between acting for the right reasons and being creditworthy. One way to flesh this is out is *Link*:

> *Link*:
> When A is required to ϕ by the members of some set of reasons S, A is creditworthy for ϕ-ing just in case A ϕs because of the members of some subset of S that are sufficiently strong to require A to ϕ.[58,59]

[57] This type of view is developed in the case of actions in Wedgwood (2006, 2007). I further develop this thought in Lord (forthcoming c) and Lord (MS, ch. 5).

[58] Arpaly (2003) and Markovits (2010) defend a very similar view to this about moral worth. If you replace 'creditworthy' with 'moral worth', then you get something very close to Arpaly's and Markovits's view. It also resembles, of course, an old Kantian idea.

[59] The bit about subsets of S is necessary, because in some cases there are multiple proper subsets of the set of all the reasons to ϕ that require one to ϕ. At least in some of these cases, it seems one can be creditworthy in ϕ-ing just so long as one acts for the reasons in one of those subsets.

When you combine *Link* with the claim that I don't act for the right reasons when I buy Anne the hat because it will make me happy, you get a principled explanation for my not being creditworthy.

With *Link* in hand, let's return to the objectivist. To repeat, the first premiss isn't a problem for the objectivist either if it is false or if, my arguments before notwithstanding, it's possible to act for the right reasons when you are ignorant of them. Let's examine each possibility in turn, starting with the latter.

One type of objectivist will insist that when you are required to ϕ because of some set of reasons *S*, you can always ϕ because of the members of *S*. There are two different kinds of cases that are relevant. *Day's End* is an instance of the first type of case. In that type of case, not only are you not in a position to know the relevant fact, you don't even believe it. I take it that it's extremely plausible that, in this type of case, you can't ϕ because of the reason. You can't refrain from flipping the switch because it will badly burn your neighbour if you don't even believe that your neighbour will be badly burned if you flip the switch. For reasons that will become clear, I don't think this yet decisively tells against this first strategy.

The second kind of case is like *Delusional Andy* (and the extension of *Day's End* where you irrationally believe flipping the switch will burn your neighbour). In that case, you do believe the relevant things, though irrationally. This first objectivist reply has serious hope of sticking if it's possible to do the right thing for the right reason in these cases. This is because it's always possible in the relevant sense of possibility to go from the first type of case to the second. Recall that something is possible in the relevant sense only if nothing about one's epistemic position bars one from doing that thing. Although one's epistemic position does bar one from knowing the relevant fact in both kinds of case, nothing about one's epistemic position in either type of case bars one from believing the relevant proposition.

Unfortunately for the objectivist, it's simply not plausible that you can do the right thing for the right reason in the second type of case. We can see this by reflecting on *Link*. It's very plausible that Delusional Andy is *not* creditworthy for his moving out and filing for divorce, whereas Surprised Andy is. Similarly, it seems as if you are not creditworthy in the extension of *Day's End* where you refrain from flipping the switch because you irrationally believe that flipping the switch will badly burn your neighbour. It might be that you end up doing what's best. But you just got lucky, and hence are not creditworthy. It follows from *Link* that Delusional Andy doesn't move out

for the right reason, nor do you refrain from flipping the switch for the right reason. If this is right, then believing the relevant fact (and acting in light of that belief) is not sufficient for acting for the right reason. Thus this first strategy fails.

The second route the objectivist can take is to deny (1). I think *Link* sheds light on the plausibility of this, as well. If you deny (1), then you are committed to thinking that there are cases where (i) the members of some set of reasons *S* make it the case that you ought to ϕ, and (ii) it's impossible to ϕ because of the members of *S*. It follows from *Link* that there are cases where the members of *S* make it the case that you ought to ϕ and it's impossible for you to ϕ and be creditworthy for ϕ-ing. I think this is deeply implausible.

It's implausible because it robs our full-stop obligations of a certain kind of *action-guidingness*. If we cannot perform the actions we are obligated to perform in a way that deserves credit, then those obligations are not action-guiding. To put it another way, if our full stop obligations are action-guiding, then our actions can be guided by the facts that determine our full-stop obligations. The paradigm of being guided by some fact is acting because of that fact. But we've just seen that in denying (1), the objectivist is denying that we can always act because of the right-makers of the acts that we are obligated to perform. Thus, there is an obvious way in which the objectivist has to deny that our obligations are always action-guiding. This seems very implausible.

The point can be put less abstractly. Consider *Day's End* again. If objectivism is true and (1) is false, then there is simply no way you can do what you ought to for the right reason. In other words, the only way you can do what you ought to do is by being an idiot. You'll have to do *something* stupid in order to do what you ought to. There's no way for the right-makers of your act to get any legitimate grip on you. But, we're supposed to believe, they require you stumble around in the dark all the same. Again, this seems implausible.

Thus, it looks as if (1) is enough to topple the objectivist. This isn't enough for me, however. This is because I not only want to defeat the objectivist, I want to establish *Possessed Reasons*. So we need to consider (2), as well. (2) holds that possessing *r* is a necessary condition for ϕ-ing because of *r*.

I think the most plausible route for denying (2) is to hold that in order to possess *r*, you have to stand in some positive justificatory relationship to *r*—for instance, you have to know *r*—but one needn't stand in this relation to *r* to ϕ because of *r*. The most natural version

of this view holds that you merely have to believe *r* in order to ϕ because of *r*.[60] Fortunately for us, we needn't dwell too much on this proposal, for we have already seen that it is implausible. It's implausible because it's implausible that merely believing *r* (and acting in light of this belief) is sufficient for ϕ-ing because of *r*. If this were true, then it would follow from *Link* that Delusional Andy is creditworthy for moving out. But he's not.[61] Thus, it seems as if you have to stand in some positive justificatory relationship to *r* in order to ϕ because of *r*. If you also have to stand in some positive justificatory relationship to *r* to possess *r*, it's plausible that in order to ϕ because of *r*, you must possess *r*.

Even if this is right, I've only established that a necessary condition for a set of reasons *S* to make it the case that you ought to ϕ is that you possess the members of *S*. This is significant in the dialectic with the objectivist, for it entails that objectivism is false. But it's not a full vindication of *Possessed Reasons*. In the rest of the section I'll argue for two claims that get us closer to a vindication of *Possessed Reasons*.

The first claim is a strengthening of the necessity condition. Not only has it to be that you possess the members of *S*, it has to be that the members of *S* are *decisive reasons* to ϕ. If they are decisive, then (i) they aren't defeated by other reasons you possess—that is, the members of *S* will be weightier than the reasons you possess not to ϕ—and (ii) they are much weightier than the reasons you possess for any alternative action.[62] Thus, I think we should replace the conclusion of the above argument with *Necessary*:

> *Necessary*:
> If the members of some set of reasons *S* make it the case that you ought to ϕ, then the members of *S* are possessed decisive reasons to ϕ.

[60] It's possible to hold that in order to possess *r* you have to know *r*, but in order to ϕ because of *r* you just have to justifiably believe *r*. I can't see any motivation for this, though. If you are going to think that there are two different epistemic relations involved, it's more plausible to think that one has to stand in some positive epistemic relation (such as knowledge or justified belief) in order to possess, but one doesn't need to stand in a positive epistemic relation in order to have the ability to act for a reason.

[61] This argument works just as well against views that accept (2) but hold that mere belief is sufficient for possession (for example, Schroeder 2008).

[62] I add the second condition because there are some cases—call them *sweetening cases*—where one has weightier reason to ϕ than ψ, but it seems as if it is still permissible to ψ. Here's one such case. Suppose I am choosing between two different job offers. They are both good jobs and I have roughly equal reason to take either. Now suppose that one job offers me a $50 signing bonus. I thus come to have slightly weightier reasons to take that job. But it seems wrong to me (and to others) to think that this means I ought to take that job. Taking the other job still seems permitted (not all agree with this; see Hare (2010)).

I don't think this strengthened necessity condition needs much more defence. It's very plausible that a set of reasons needs to decisively support ϕ-ing in order to make it the case that you ought to ϕ. If the members of S are defeated or are not much weightier than the alternatives, then it's independently plausible that they don't make it the case that you ought to ϕ.

We'd be much closer to a vindication of *Possessed Reasons* if we were to show that *Sufficient* is true:

> *Sufficient*:
> If the members of S are decisive reasons to ϕ that you possess, then the members of S make it the case that you ought to ϕ.

If *Necessary* is true and *Sufficient* is false, then there is some condition over and above possessing decisive reasons to ϕ that must be met in order for it to be the case that you ought to ϕ. Obviously there are many possible conditions that one could propose. I don't have a good enough imagination to think of all of them (nor do I have the space to consider them all). But it's worth considering two of them.

The first might be some kind of ability condition. It might be that 'ought' implies 'can', but that 'having decisive reasons' doesn't imply 'can.' I agree (obviously!) that there is some ability condition. Not only do I think that you must be able to ϕ in order to be required to ϕ, I think you must be able to ϕ because of the right-makers in order to be required to ϕ. But whatever plausibility these claims have when it comes to what you ought to do seems to apply equally well to what you possess decisive reasons to do. Any example that provides evidence for 'ought' implies 'can' will, I conjecture, provide equally good evidence for 'possess decisive reason' implies 'can'. A plausible generic claim in the vicinity is something like: if r is a reason to ϕ, then it's possible to ϕ because of r.

A second proposal is some type of hypological condition. For example, perhaps it's true that if you ought to ϕ, then you are blameworthy for not ϕ-ing, but it's not the case that you are always blameworthy for not doing what you possess decisive reasons to do. If this were true, then there would be cases where you possess decisive reasons to ϕ, but wouldn't be blameworthy if you didn't ϕ, and thus it wouldn't be the case that you ought to ϕ. One problem with this proposal is that the hypological condition on 'ought' is false. It's just not true that you are always blameworthy for not doing what you ought to. There are simple recipes for counterexamples. One is

to take a case where your pet theory of obligation says you ought to ϕ at time t. At t_1, bombard your attention with useless data (think: loud noises or bright flashing lights). You can set the case up so that it is very hard to do what you ought to (and thus most of the time you won't), but nevertheless, the conditions that obligate you to ϕ are still in place. In this type of case, you will seem blameless for not ϕ-ing, because of how hard ϕ-ing becomes, even though you will still be obligated to ϕ.

Furthermore, even if you think that you are always blameworthy for not doing what you ought to, it's hard to see why you would want to deny that you are always blameworthy for not doing what you possess decisive reasons to do. The most plausible reason I can think of is that what you ought to do is transparent in some way, but what you possess decisive reason to do isn't (perhaps because being decisive isn't transparent). You could think this, but I'm not sure why you would. If what you ought to do is transparent, then it seems as if it will always be transparent what the facts support, but this is just another way of saying that it will be transparent what the reasons you possess decisively support.

These two proposals are not plausible, and thus it seems that having decisive reasons to ϕ is sufficient for it being the case that you ought to ϕ. This gives us *Necessary and Sufficient*:

> *Necessary and Sufficient*:
> The members of S are decisive reasons to ϕ that you possess iff the members of S make it the case that you ought to ϕ.

Necessary and Sufficient is, I think, just a precisification of *Possessed Reasons*. It tells you how your obligations are a function of the reasons you possess.

5. Summary of results (or why you ought to be rational)

As we saw at the beginning, rationality has traditionally been seen to be on steady deontic footing. Finding out you would be irrational for ϕ-ing has struck most philosophers and laypeople alike as bearing very directly on what ought to be done. Despite this, recent work in metaethics makes it surprisingly plausible that rationality has no direct effect on what ought to be done.

As we saw in §2.2, this is largely because it has become popular to think that rationality is constitutively tied to coherence. This

connection has been spelled out in two different ways in the literature. According to the narrow-scope view, you are required by rationality to have certain attitudes when you have certain other attitudes. For example, you are required to intend to ψ when you intend to ϕ and believe that in order to ϕ you must ψ. According to the wide-scope view, you are never required to have particular attitudes. Rather, you are directly required to be coherent—you are required to [intend to ψ if you intend to ϕ and believe that in order to ϕ you must ψ].

The problem is that it is very hard to show that the requirements posited by these two views have deontic significance. The narrow-scope view seems hopeless, because it allows for a pernicious kind of bootstrapping. It is very implausible that one ought to or has reason to intend to ψ*simply because* one intends to ϕ and believes that in order to ϕ one must intend to ψ. On the other hand, it is mysterious how it could be the case that the wide-scope requirements are deontically significant. This is because it is mysterious what reasons there could be to directly be coherent. We might have reasons to have particular attitudes, and they might weigh up such that we always have sufficient reason to have attitudes that cohere. But this is not enough to vindicate the deontic significance of the wide-scope requirements.

These results have led many to scepticism about the deontic significance of rationality. This, I think, is a mistake. We should question one of the starting points—namely, that rationality is constitutively tied to coherence. Instead, I think we should think that rationality consists in correctly responding to the objective reasons one possesses. At the beginning of §3 I dubbed this claim Reasons Responsiveness.

It was established in §3.2 that if Reasons Responsiveness is true, then rationality is weakly deontically significant. This is because there is always reason to do what you possess reason to do. So if you always possess reasons to do what rationality requires, then there will always be reasons to do what rationality requires. This only goes so far, though. It would still be a considerable win for sceptics about the deontic significance of rationality if rationality turns out to only be weakly deontically significant. In order to fully vindicate the deontic significance of rationality, one needs to show that it's strongly significant—that is, that you ought to be rational.

The last section aimed to show this indirectly. The main claim of the last section was that what you ought to do is determined by the reasons that you possess. The plausibility of this claim is independent

of any considerations having to do with rationality—this is why it was indirect.

When you combine §3 with the results of §4, you see that Reasons Responsiveness can vindicate the strong deontic significance of rationality. If what you are rationally required to do is what you possess decisive reasons to do, and what you ought to do is what you possess decisive reasons to do, then the requirements of rationality *just are* the requirements you ought to comply with. Indeed, on my picture, there is no space between what the (possessed) reasons require, what rationality requires, and what you ought to do. For this reason, I think we should drop the Broomian idea that rationality has a *source* that is independent of normative reasons.[63] It is not the case that reasons and rationality are independent of each other in the way Broome envisions. If they were, then the best we could hope for is a correlation between what rationality requires and what you ought to do. I believe this paper vindicates something stronger: what you are rationally required to do just is what you ought to do.

It is a serious virtue of my view that it can give such a plausible account of why rationality is strongly deontically significant. This is especially weighty, since it's completely mysterious how rival views can even account for the weak deontic significance of rationality, let alone the strong deontic significance. Given that it is a truism that you should be rational, the fact that only my view can vindicate the strong deontic significance of rationality is a strong reason to accept my view—one which you now possess.[64]

References

Arpaly, Nomy 2003: *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.

[63] As a referee pointed out to me, this makes it seem as if rationality is a source of normative reasons that competes with the other sources. This is not a natural view. Fortunately my view allows us to avoid it.

[64] Thanks to audiences at Penn, UCSB, Franklin and Marshall, and SUNY Buffalo for discussion of this paper, and audiences at Princeton, Edinburgh, MadMeta and Georgetown for discussion of many of the main ideas. Thanks also to Tom Kelly, Michael Smith, Gideon Rosen, Andrew Huddleston, Liz Harman, Jonathan Way, Daniel Whiting, Conor McHugh, Rishi Joshi, Anthony Cross, Eden Lin, Nat Tabris, Simon Cullen, Daniel Fogal, Richard Yetter Chappell, Shyam Nair, Alex Worsnip, Julien Dutant, Karl Schafer, Susanne Mantel, and several anonymous referees for helpful discussion and written comments. My biggest debt is to Kurt Sylvan. I thank him for extremely generous feedback over the course of many years.

Bedke, Matthew S. 2009: 'The Iffiest Oughts: A Guise of Reasons Account of End-Given Reasons and End-Given Oughts'. *Ethics*, 119(4), pp. 672–98.
Björnsson, Gunnar, and Stephen Finlay 2010: 'Metaethical Contextualism Defended'. *Ethics*, 121(1), pp. 7–36.
Bratman, Michael E. 2009a: 'Intention, Belief, Practical, Theoretical'. In Simon Robertson (ed.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
—— 2009b: 'Intention, Practical Rationality, and Self-Governance'. *Ethics*, 119(3), pp. 411–43.
Broome, John 1999: 'Normative Requirements'. *Ratio*, 12, pp. 389–419.
—— 2005a: 'Does Rationality Give Us Reasons?' *Philosophical Issues*, 15, pp. 321–37.
—— 2005b: 'Have We Reason to Do as Rationality Requires? A Comment on Raz'. *Journal of Ethics and Social Philosophy*, 1, pp. 1–8.
—— 2007a: 'Does Rationality Consist in Responding Correctly to Reasons?' *Journal of Moral Philosophy*, 4, pp. 349–74.
—— 2007b: 'Requirements'. In Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson and Dan Egonsson (eds.), *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Online tribute. http://www.fil.lu.se/hommageawlodek/.
—— 2008a: 'Is Rationality Normative?' *Disputatio*, 11, pp. 153–71.
—— 2008b: 'Reply to Southwood, Kearns and Star, and Cullity'. *Ethics*, 119(1), pp. 96–108.
—— 2013: *Rationality Through Reasoning*. Oxford: Wiley Blackwell.
Burgess, Alexis, and David Plunkett 2013a: 'Conceptual Ethics I'. *Philosophy Compass*, 8(12), pp. 1091–1101.
—— 2013b: 'Conceptual Ethics II'. *Philosophy Compass*, 8(12), pp. 1102–10.
Buss, Sarah MS: 'Norms of Rationality and the Superficial Unity of the Mind'. Unpublished manuscript.
Dancy, Jonathan 2000: *Practical Reality*. Oxford: Oxford University Press.
Darwall, Stephen L. 1983: *Impartial Reason*. Ithaca, NY: Cornell University Press.
Foot, Philippa 1972: 'Morality as a System of Hypothetical Imperatives'. *Philosophical Review*, 81(3), pp. 305–16.
Gibbons, John 2001: 'Knowledge in Action'. *Philosophy and Phenomenological Research*, 62(3), pp. 579–600.

Graham, Peter A. 2010: 'In Defense of Objectivism about Moral Obligation'. *Ethics*, 121(1), pp. 88–115.
Greenspan, P. S. 1975: 'Conditional Oughts and Hypothetical Imperatives'. *Journal of Philosophy*, 72(10), pp. 259–76.
Hare, Caspar 2010: 'Take the Sugar'. *Analysis*, 70(2), pp. 237–47.
Harman, Gilbert 1985: 'Is There a Single True Morality?' In David Copp and David Zimmerman (eds.), *Morality, Reason, and Truth*, pp. 27–48. Totowa, NJ: Rowman and Littlefield. Reprinted in his *Explaining Value and Other Essays in Moral Philosophy*, pp. 77–99. Oxford: Oxford University Press, 2000.
Hill, Thomas E., Jr. 1973: 'The Hypothetical Imperative'. *Philosophical Review*, 82(4), pp. 429–50.
Hornsby, Jennifer 2008: 'A Disjunctive Conception of Acting for Reasons'. In Adrian Haddock and Fiona Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.
Hussain, Nadeem J. Z. MS: 'The Requirements of Rationality'. Unpublished manuscript.
Hyman, John 2006: 'Knowledge and Evidence'. *Mind*, 115, pp. 891–916.
Jackson, Frank 1991: 'Decision-Theoretic Consequentialism and the Nearest and Dearest Objection'. *Ethics*, 101(3), pp. 461–82.
Kavka, Gregory S. 1983: 'The Toxin Puzzle'. *Analysis*, 43(1), pp. 33–6.
Kiesewetter, Benjamin 2011: '"Ought" and the Perspective of the Agent'. *Journal of Ethics and Social Philosophy*, 5(3). http://www.jesp.org/PDF/c971_BenjaminKiesewetter.pdf.
——— 2015: 'Instrumental Normativity: In Defense of the Transmission Principle'. *Ethics*, 125(4), pp. 921–46.
Kolodny, Niko 2005: 'Why Be Rational?' *Mind*, 114, pp. 509–63.
——— 2007: 'How Does Coherence Matter?' *Proceedings of the Aristotelian Society*, 107, pp. 229–63.
——— 2008a: 'The Myth of Practical Consistency'. *European Journal of Philosophy*, 16(3), pp. 366–402.
——— 2008b: 'Why Be Disposed to Be Coherent?' *Ethics*, 118(3), pp. 437–63.
Kolodny, Niko and John MacFarlane 2010: 'Ifs and Oughts'. *Journal of Philosophy*, 107(3), pp. 115–43.
Korsgaard, Christine M. 1996: *The Sources of Normativity*. Cambridge: Cambridge University Press.
Lord, Errol 2010: 'Having Reasons and The Factoring Account'. *Philosophical Studies*, 149(3), pp. 283–96.

——— 2014a: 'The Coherent and the Rational'. *Analytic Philosophy*, 55(2), pp. 151–75.
——— 2014b: 'The Real Symmetry Problem(s) for Wide-Scope Accounts of Rationality'. *Philosophical Studies*.
——— 2015: 'Acting for the Right Reasons, Abilities, and Obligation'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 10, pp. 26–52. Oxford: Oxford University Press.
——— forthcoming a: 'Defeating the Externalist's Demons'. In Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon: New Essays on Knowledge, Justification, and Rationality*. Oxford: Oxford University Press.
——— forthcoming b: 'The Explanatory Problem for Cognitivism about Practical Reason'. In Conor McHugh, Jonathan Way and Daniel Whiting (eds.), *Normativity: Epistemic and Practical*. Oxford: Oxford University Press.
——— forthcoming c: 'On the Intellectual Conditions for Responsibility: Acting for the Right Reasons, Conceptualization, and Credit'. *Philosophy and Phenomenological Research*.
Lord, Errol MS: '*The Importance of Being Rational*'. Unpublished manuscript.
Lord, Errol and Barry Maguire 2016: 'An Opinionated Guide to the Weight of Reasons'. In Errol Lord and Barry Maguire (eds.), *Weighing Reasons*. Oxford: Oxford University Press.
Lord, Errol and Kurt Sylvan MS: 'Reasons: Wrong, Right, Normative, Fundamental'. Unpublished manuscript.
McKeever, Sean, and Michael Ridge 2012: 'Elusive Reasons'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 7, pp. 110–37. Oxford: Oxford University Press.
McNaughton, David, and Piers Rawling 2011: 'The Making/Evidential Reason Distinction'. *Analysis*, 71(1), pp. 100–2.
Marcus, Eric 2012: *Rational Causation*. Cambridge, MA: Harvard University Press.
Markovits, Julia 2010: 'Acting for the Right Reasons'. *Philosophical Review*, 119(2), pp. 201–42.
Moore, G. E. 1912: *Ethics*. Oxford: Oxford University Press.
Parfit, Derek 1997: 'Reasons and Motivation'. *Proceedings of the Aristotelian Society Supplementary Volume* 71, pp. 99–130.
Plunkett, David 2015: 'Which Concepts Should We Use? Metalinguistic Negotiations and the Methodology of Philosophy'. *Inquiry*, 58(7–8), pp. 828–74.

Plunkett, David and Timothy Sundell 2013: 'Disagreement and the Semantics of Normative and Evaluative Terms'. *Philosophers' Imprint*, 13(23).
Raz, Joseph 2005: 'The Myth of Instrumental Rationality'. *Journal of Ethics and Social Philosophy*, 1.
Ross, Alf 1941: 'Imperatives and Logic'. *Theoria*, 7, pp. 53–71.
Ross, Jacob 2006: *Acceptance and Practical Reason*. Ph. D. thesis, Rutgers University.
—— 2012: 'Rationality, Normativity, and Commitment'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 7, pp. 131–88. Oxford: Oxford University Press.
Scanlon, T. M. 2001: 'Thomson on Self-Defense'. In Alex Byrne, Robert Stalnaker and Ralph Wedgwood (eds.), *Fact and Value: Essays on Ethics and Metaphysics for Judith Jarvis Thomson*, pp. 199–213. Cambridge, MA: MIT Press.
—— 2008: *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
Schroeder, Mark 2004: 'The Scope of Instrumental Reason'. *Philosophical Perspectives*, 18: *Ethics*, pp. 337–64.
—— 2005a: 'The Hypothetical Imperative?' *Australasian Journal of Philosophy*, 83(3), pp. 357–72.
—— 2005b: 'Instrumental Mythology'. *Journal of Ethics and Social Philosophy*, 1.
—— 2008: 'Having Reasons'. *Philosophical Studies*, 139(1), pp. 57–71.
—— 2009a: 'Buck-Passers' Negative Thesis'. *Philosophical Explorations*, 12(3): Practical Reasoning and Normativity, pp. 341–7.
—— 2009b: 'Means–End Coherence, Stringency, and Subjective Reasons'. *Philosophical Studies*, 143(2), pp. 223–48.
—— 2010: 'Value and the Right Kind of Reason'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 5, pp. 25–55. Oxford University Press.
—— 2015a: 'Hypothetical Imperatives: Scope and Jurisdiction'. In Mark Timmons and Robert N. Johnson (eds.), *Reason, Value, and Respect: Kantian Themes from the Philosophy of Thomas E. Hill, Jr.*, pp. 89–100. Oxford University Press.
—— 2015b: 'What Makes Reasons Sufficient?' *American Philosophical Quarterly*, 52(2), pp.159–70.
Smith, Matthew Noah 2016: 'One Dogma of Philosophy of Action'. *Philosophical Studies*, 173(8), pp. 2249–66.
Smith, Michael 1994: *The Moral Problem*. Oxford: Blackwell.

——— 2008: 'Consequentialism and the Nearest and Dearest Objection'. In Ian Ravenscroft (ed.), *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*, pp. 237–66. Oxford: Clarendon Press.
Southwood, Nicholas 2008: 'Vindicating the Normativity of Rationality'. *Ethics*, 118(1), pp. 9–30.
Thomson, Judith Jarvis 1990: *The Realm of Rights*. Cambridge, MA: Harvard University Press.
——— 2008: *Normativity*. Chicago: Open Court Press.
Way, Jonathan 2010a: 'Defending the Wide-Scope Account of Instrumental Reason'. *Philosophical Studies*, 147(2), pp. 213–33.
——— 2010b: 'The Normativity of Rationality'. *Philosophy Compass*, 5(12), pp. 1057–68.
——— 2012: 'Explaining the Instrumental Principle'. *Australasian Journal of Philosophy*, 90(3), pp. 487–506.
——— 2013: 'Intentions, Akrasia, and Mere Permissibility'. *Organon F*, 20(4), pp. 588–611.
Wedgwood, Ralph 2006: 'The Normative Force of Reasoning'. *Noûs*, 40(4), pp. 660–86.
——— 2007: *The Nature of Normativity*. Oxford: Oxford University Press.
Williams, Bernard 1980: 'Internal and External Reasons'. In Ross Harrison (ed.), *Rational Action: Studies in Philosophy and Social Science*, pp. 17–28. Cambridge: Cambridge University Press. Reprinted in his *Moral Luck: Philosophical Papers* 1973–1980, pp. 101–13. Cambridge: Cambridge University Press, 1981.
Zimmerman, Michael J. 2008: *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.

