**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2019

A reason to be rational

Mildenberger, Carl David

DOI: <https://doi.org/10.1080/0020174x.2018.1470570>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-191832>

Journal Article

Originally published at:

Mildenberger, Carl David (2019). A reason to be rational. *Inquiry*, 62(9-10):1008-1032.

DOI: <https://doi.org/10.1080/0020174x.2018.1470570>

A reason to be rational

Abstract. This essay argues that in spite of the powerful arguments by Kolodny and Broome there *is* a reason to be rational. The suggested reason to be rational is that if an agent complies with rational requirements the people around him, as well as he himself, will be able to explain and predict his attitudes. Rationality allows us to make sense of an agent's attitudes in terms of his other attitudes. This form of explainability is valuable, because it provides us with greater comprehension as regards an agent's attitudes. Thus, I argue that there is an instrumental reason to be rational.

Keywords: rationality; reason; normativity; explanation; wide-scope requirements

1 Introduction

In recent years, notably John Broome (e.g. 1999, 2007a, 2013) has advocated a conception of rationality as the coherent relation of an agent's attitudes. According to this conception, we can think of rationality as a set of requirements that govern the internal relations of an agent's attitudes. That is, rational requirements are not about how an agent's attitudes relate to facts of the external world, but only about how an agent's attitudes are related amongst each other. Take the example of the following requirement (Broome 2007a, 162).

Modus ponens (MP): Rationality requires of N that, if N believes p and N believes that if p then q , and if it matters to N whether q , then N believes q .

MP states how an agent's attitudes have to be related in order to fulfill the requirement: they have to be coherent. The agent must not hold beliefs that cannot all be true together.¹

There is considerable controversy over whether rational requirements are *wide-* or *narrow-scope*. Broome (2007b) argues in favor of wide-scope requirements. That is to say that rationality governs the relationship of an agent's attitudes rather than to require the agent to hold any specific, singular attitude. Broome's example of MP does not tell the agent to hold the belief that q . Instead, a bit simplified and with added brackets for clarification, it requires of the agent that (if he believes that p , and also believes that if p then q , then he believes that q). The requirements of rationality govern the entire conditional of MP, not only its consequent.² There also is considerable controversy over whether rational requirements are *state* or *process requirements*; where state requirements are concerned with attitudes we contemporaneously hold, and process requirements tell us which attitudes to form, revise, or drop over time (Kolodny 2005, 515–18).³

However, in spite of these controversies, there is widespread agreement that if we conceive of rationality in terms of requirements, we seem to face a peculiar problem. Broome (2013, 192–205), for one, acknowledges that on his account it looks like we have no reason to comply with rational requirements. He cannot provide us with a positive answer to the normative question of whether,

¹ For requirements of theoretical rationality, the coherence called for is straightforward logical coherence. In the case of practical requirements involving both beliefs and intentions, what kind of coherence is called for is less clear (cf. Reisner 2011, 51–53).

² Kolodny (2005, 2007) is the most prominent defender of the narrow-scope view. For a defense of wide-scope requirements of rationality – in light of Schroeder's (2004) important criticism that wide-scope requirements do not capture the *asymmetry* of rational requirements – see e.g. Brunero (2010) and Way (2011).

³ Kolodny (2007) thinks that at least some rational requirements are process requirements, whereas Broome (2007b) is skeptical, and Reisner (2009) even more so.

if rationality requires us to F , we have a reason to F . Broome is skeptical about whether rationality is normative, we seemingly have no reason to be rational. Kolodny (2005, 542–47) reaches a similar conclusion.⁴

The purpose of this paper is to show that there is a reason to be rational, where rationality is understood as a set of wide-scope requirements.⁵ This reason is that, if an agent complies with rational requirements, the people around him, as well as he himself, are able to explain and predict his beliefs and intentions. Rationality allows us to make sense of an agent's attitudes in terms of his other attitudes. This form of being explainable is valuable, because it provides us with greater comprehension as regards an agent's beliefs and intentions. It allows us to understand why agents believe that p or intend to F . My suggested reason to be rational is that being rational entails having explainable and predictable attitudes, which is associated to clear benefits. Thus, I argue that there is an instrumental reason to be rational.⁶

I shall first briefly outline desiderata a reason to be rational has to fulfill (Section 2). I will then argue how being rational makes us explainable and predictable to ourselves and others (Section 3), and how this is reason-giving (Section 4). Section 5 shows how my suggested reason to be rational is able to meet the desiderata and deals with some objections. Notably, I will address the worries that it is not always beneficial to be explainable and predictable and that we can also become explainable and predictable by ways other than being rational.

⁴ I will have a closer look on why Broome and Kolodny think so in Section 2.

⁵ I focus on wide-scope requirements, as it seems to be virtually unanimously agreed that narrow-scope requirements could not be normative for the threat of bootstrapping (e.g. Kolodny 2005; Broome 2007b).

⁶ Broome remarks that identifying an instrumental reason to be rational only means to show that rationality is *weakly* normative. We may speak of weak normativity in case that '[n]ecessarily, if rationality requires N to F , there is a reason for N to F ' (2007a, 163–64). Arguing in favor of rationality being weakly normative, as I do here, does not help to establish the claim that rationality *itself* is a source of normativity (Broome 2007a, 165). It only associates rationality with normativity in that rationality is always able to borrow normativity from somewhere else.

2 Desiderata of a reason to be rational

Testing different candidates for reasons to be rational, Kolodny (2005, 542–47) and Broome (2007a, 171–77) compile an impressive list of considerations in favor of being rational that, while promising *prima facie*, do not withstand scrutiny. In the course of their analysis, Kolodny and Broome establish the following desiderata a genuine reason to be rational has to meet.⁷

First, a reason to be rational has to be *substantive*, in the sense that it features clear benefits that arise from being rational. If, for example, you were more likely to believe and do what you ought as a result of being rational, then this would constitute a substantive reason to be rational (Kolodny 2005, 543).⁸

Second, a reason to be rational has to be *general*, i.e. it has to exist ‘for any believer or intender, no matter what his circumstances’ (Kolodny 2005, 543). Kolodny’s candidate reason of being more likely to believe and do what you ought, for instance, is not general. Its generality is undermined by *error cases*. Kolodny highlights that we might hold a perfectly coherent set of beliefs $b_1, b_2, b_3, \dots, b_n$ with the beliefs all being false. In this case, he argues, the set of coherent beliefs is not what we ought to believe. Thus, this candidate reason might be said to be substantive but not sufficiently general (Kolodny 2005, 543).

Acknowledging that error cases always are a threat, Broome (2007a, 172) believes that there cannot be a reason to satisfy individual requirements of rationality *directly*. But this leaves open the possibility of there being an *indirect* reason to be rational. There might be something that we ought to have or be which, in turn, makes us comply with rational requirements. Broome (2007a, 172,

⁷ I will only consider desiderata *instrumental* reasons to be rational have to meet according to Broome and Kolodny, thus leaving aside desiderata for potentially existing *final* or *intrinsic* reasons to be rational. This seems appropriate since I shall myself put forward an argument for an instrumental reason to be rational.

⁸ Kolodny is assuming here a rather uncomplicated relationship between intentions and actions, i.e. that we just do whatever we intend. I follow Kolodny and others in not problematizing this step from intention to action.

176) examines the possibilities that we might ought to have the ‘rational faculty’ (a bundle of abilities and dispositions which *causes* us to satisfy rational requirements), or ought to be ‘fully rational’ (which *entails* satisfying all rational requirements).

Such indirect reasons to be rational feature additional desiderata. First, a reason to be rational must not rely on a dubious *principle of inference* which would allow normativity to be transmitted to causal effects – as it is the case for the rational faculty. Broome (2007a, 175) gives the example of the side effects of a drug. Just because you have a reason to take the drug (e.g. the drug will cure you), this does not mean that you have a normative reason to have the side effects caused by the drug.

Second, a reason to be rational must not rely on a *psychological impossibility* – as it is the case for the property of being fully rational (Broome 2007a, 177). Given our cognitive limits nobody can be as rational as to fulfill all rational requirements. As being fully rational is not even psychologically possible for us, it can hardly support a normative reason to be rational.

These are then, the four suggested desiderata a reason to be rational has to meet. The reason to be rational has to be substantive and general, and it must not rely on a dubious principle of inference or a psychological impossibility.

3 Rationality and traceability

In their search for a reason to be rational both Broome and Kolodny seem to have overlooked a promising candidate for an instrumental reason. The general idea which underlies this new candidate reason is the following. What we are looking for when we are looking for an instrumental reason to be rational (where being rational is understood as complying with wide-scope requirements) is some kind of benefit that arises out of our attitudes complying with certain *formal* requirements. We are looking for a reason for complying with a certain formal structure that is entirely independent of the *contents* of our attitudes.

Now, there is something which can be achieved because of the formal structure our attitudes comply with when we are rational. Namely, if we are complying with rational requirements, this makes our attitudes *traceable* for us and for others around us. My central claim is that being rational makes the agent's attitudes *traceable* for herself and those around her, which is associated to clear benefits.

Let me explain what I mean by 'being traceable' before I address why it is valuable. A person A 's (as in *agent*) attitudes are traceable in a specific situation, if A herself or B (as in *bystander*) can *explain* why A believed or intended something, given that A 's other attitudes are known. Traceability also means that A or B are to a certain extent able to *predict* what A believes or intends, given that A 's other attitudes are known.

This is quite a mouthful. But consider a simple example to intuitively grasp what traceability amounts to. Suppose B wants to explain why A believed that q . If B knows that A believed that p , and that she also believed that if p then q , and if A 's attitudes are *traceable*, then B can explain why A believed that q . In this example, having coherent attitudes, i.e. conforming to rational requirements, is what makes one's actions traceable. For a rational agent we can explain an attitude she held or currently holds in terms of her other attitudes. If we know that an agent believes that p , and if we also know that she believes that if p then q , then we can explain her belief that q in terms of these two other attitudes. Her belief that q can be explained by pointing to the fact that – in accordance with the coherence requirements of rationality (and notably MP) – if she believes that p , and that if p then q , she also believes that q .

There are important parallels here to Hempel and Oppenheim's (1948) deductive-nomological account of scientific explanation. Hempel and Oppenheim essentially hold that in an explanation the *explanandum* will be the conclusion of a valid argument, the premises of which will contain general laws and initial conditions. Only if there is a general law covering a certain phenomenon will we be able to *trace* (to use my term) how we proceed from initial conditions to the *explanandum*. With respect to our context, I suggest that the *explanandum* is a specific belief or intention of the

agent, while the rational requirements can assume the role of the general law, and the agent's other attitudes are the initial conditions.⁹ If people behave rationally, then the rational requirements they conform to may serve us as the general law we need in order to trace their beliefs and intentions in terms of their other attitudes.

In our example involving A and B , B is able to give a deductive-nomological explanation of why A believes that q . At this point, it is easy to see that it is not the knowledge of an agent's initial attitudes alone that lets us explain her attitudes in terms of her other attitudes. The abundance by rational requirements contributes decisively to the explainability, because we need a general law that links the attitude we are explaining to the other attitudes in terms of which we are explaining. My suggested instrumental reason to be rational precisely is that if an agent is rational, she and the people around her who interact with her are able to trace her beliefs and intentions.

The example just given involving A and B is an example of how, in a *social condition*, i.e. when people are interacting with each other in some way, a bystander is able to *explain* an agent's attitudes. The other aspect of traceability in the social condition is that it enables bystanders to *predict* specific attitudes of a rational agent to a certain extent.

Suppose that an agent intends to G and also believes that intending to F is a necessary means to G ing. What can we predict about further attitudes of hers? If the agent is rational, then her attitudes conform to the following requirement (Broome 2007a, 162).

Means-end (ME): Rationality requires of you that, if you intend to G , and if you believe your F ing is a necessary means to your G ing, and if you believe you will not F unless you intend to F , then you intend to F .

⁹ For what follows I will adopt a deductive-nomological idea of tracing. I will also adopt Hempel and Oppenheim's idea of symmetry between explanation and prediction by referring to the process of tracing as 'explanation' if the *explanandum* lies in the past and as 'prediction' if it lies in the future.

If we know she has the antecedent attitudes and that she is coherent with respect to ME, then we can predict that she has the intention to F . That is something. For example, if we know that an agent intends to drive safely from Edinburgh to London, and that she believes that in order to drive safely from Edinburgh to London, she must intend to drive on the left side, then we can predict that she intends to drive on the left side. For practical purposes, i.e. for me as a driver who drives in the opposite direction, that is all I need to know for the moment when I meet her on the road. After all, based on only two attitudes, by which I mean based on very sparse knowledge, we can predict a third one. In this respect, rationality seems to have quite a considerable predictive power.¹⁰

Tracing can also be performed in the *individual condition*. An agent can trace her own attitudes. The process is analogous to that in the social condition. In the individual condition $\mathcal{A}$ is facing her own attitudes in the same way in which a bystander faces them in the social condition. For instance, in some situation, it might not be enough for $\mathcal{A}$ to be aware that she believes that q , but she might want to explain why she believes that q . If she is traceable, she can resort to the same explanation available to a bystander and say: ‘I believe that q because I’m rational, and because I believe that p , and that if p then q ’.

¹⁰ In this case, our prediction is based on coherence and on the fact that we know the agent’s initial attitudes. One might object that the agent could also stay coherent with respect to ME by, say, dropping the antecedent-intention and adopting the consequent-intention not to F at the same time. Thus, one might argue, we are not really able to predict, based on coherence and knowledge of initial attitudes alone, that the agent will adopt the intention to F . This objection crucially depends on what we mean by ‘knowing an agent’s initial attitudes’. I take it to be the most natural meaning that knowledge of an agent’s initial attitudes implies that they do not simply change in the situation for which we make our deductive-nomological prediction. In any case, this is the meaning I am adopting here. Thus, the objection does not apply.

4 The benefits of being traceable

I have argued that being rational entails being traceable. I shall now turn to elaborating what kind of benefits arise from being traceable. This is important because without being traceable having clear benefits, being traceable is unable to sustain an instrumental reason to be rational.

The general idea is that being traceable is valuable because it is a form of being explainable or predictable. Namely, it is that kind of being explainable or predictable which is concerned with deductive-nomologically explaining or predicting an agent's attitudes in terms of her other attitudes and her compliance with rational requirements. To show how my suggested reason to be rational works in detail, it is best to separately examine the instrumental benefits of traceability in four different cases: (i) predictability and (ii) explainability in the social condition, as well as (iii) predictability and (iv) explainability in the individual condition. This is because these four cases differ with respect to how substantive and how general the instrumental benefits granted by traceability are. Roughly speaking, the benefits of traceability are most substantive in the social condition, in the sense of traceability typically providing the biggest benefits there, but more generally realizable in the individual condition. I will now examine the four cases in turn and argue that ultimately a general and substantive instrumental reason to be rational needs to be based on the benefit explainability grants in the individual condition – namely, self-comprehension.

The case in which the instrumental benefits of traceability are most palpable is (i) predictability in the social condition. Being predictable is valuable in every society, because societies are all about coordinating and cooperating with each other. It is because of coordination and cooperation that those who live in a society are able to satisfy more needs than they would be able to in isolation. But successful coordination and cooperation are in need of agents with predictable actions and beliefs.¹¹ However, it is not only beneficial for other people or society as a whole that an agent is

¹¹ Bratman (1987, chap. 3) makes comparable claims about the value of interpersonal coordination and the respective value of coherent intentions and plans.

predictable. If we change perspective and ask what benefits the agent can secure from her being predictable, these benefits are a mirror image. When interacting with others, usually it is good for the agent to be predictable. For instance, you will only be trusted, assigned to important positions, and be allowed to work independently as a part of a larger group if you are predictable. It also seems plausible that we can most efficiently benefit agents, i.e. further their individual ends, if they are predictable. You benefit whenever somebody can predict your actions and thus can easily oblige you, e.g. your partner catching you outside the place you work to walk home along with you.

In contrast, the clearest case for the benefits of (ii) explainability in the social condition rests on the value of mutual comprehension and communication. It is hard to see how mutual comprehension and communication would be possible without the attitudes of other people being explainable in some form.¹² And mutual comprehension and communication are crucial fundaments of every society. If we are able to explain why some agent believes that q , this immediately opens up richer ways of interaction than if we only knew that she believes that q . For instance when we want to convince her that $\neg q$. Also, whenever we build on the knowledge of past generations, explainability of past agents' attitudes is key. Again, the benefits the agent can secure from her being explainable are a mirror image. For example, you benefit from being explainable in the social condition every time you present an argument.

¹² Compare Davidson (1973) who claims that whenever we are trying to interpret what somebody is saying, be it in a foreign language we do not speak or in English, we necessarily have to assume that there is coherence among that person's attitudes. This is because 'disagreement and agreement alike are intelligible only against a background of massive agreement' (1973, 324). According to Davidson, since we ourselves use rational requirements in order to distinguish valid combinations of beliefs, we can never explain another person's beliefs without assuming them to be coherent. Unlike Davidson, my account of tracing does not conceive of coherence as a necessary condition for making sense of an agent's attitudes (cf. Section 5). Furthermore, I am making no assumption with respect to whether the interpreting/tracing person's attitudes also are coherent.

As substantive as I take these instrumental benefits in the social condition to be, three concerns about basing a reason to be rational on them alone come to mind. First, predictability and explainability seem to be necessary conditions for things like cooperation or communication. In order to say, for example, that the benefits of cooperation and interpersonal coordination give us reason to be predictable we would have to accept transmission of reasons across necessary means – which one might consider a dubious principle. Second, one might object that it is not always beneficial for the agent herself to be traceable in the social condition. There seem to be quite some instances in which it is in your interest to sabotage rather than facilitate the effective tracing of your attitudes; e.g. if you are a famous diamond thief. Third, we simply are not always in the social condition. In view of these three concerns it becomes clear that in order to find a genuine instrumental reason to be rational we need to move on to the benefits of being traceable in the individual condition.

Intuitively, the benefits of (iii) predictability in the individual condition are substantive and significantly more generally realizable than in the social condition. It always seems a good thing to be able to predict one's own attitudes. For instance, it allows us to plan our own actions.¹³ Yet, whatever the benefits of an agent's being able to predict her own attitudes are, they do not seem to fully apply in our context. This is because of *conceptual* worries concerning the concept of prediction.

How can we predict anything in the individual condition based on complying with *state* requirements like MP? It is hard to argue that I am in any meaningful sense predicting that I believe that q , when I derive my belief that q from my other beliefs that p and that if p then q as well as my compliance with MP. Such 'prediction' neither tells me anything about the future, as I already believe that q . Nor does it tell me anything new in the vast majority of cases. For it seems that

¹³ Compare Bratman's closely related argument that the 'reason for intention coherence ... derives from a reason for self-governance' (2009, 429) – with self-governance being intrinsically valuable.

insofar as I can introspect (or otherwise directly come to know) that I believe that p and that if p then q , I could normally likewise introspect and come to know that I believe that q directly. Even more so as MP only applies to cases in which it matters to me whether q . Unlike in the social condition, in which bystanders can predict further, previously unknown attitudes of a traceable agent based on those attitudes of hers they already know, genuine prediction of this kind mostly seems to be blocked in the individual condition, if we conceive of rational requirements as state requirements. The very nature of state requirements prevents meaningful prediction on the basis of such requirements from taking place in the individual condition most of the time.¹⁴ Clearly, *process* requirements of rationality – which tell us what attitudes to form, revise, or drop going forward in time – are not affected by this conceptual worry. Genuine prediction is possible based on compliance with process requirements. But as I aim to come up with a reason to be rational that is independent of whether rational requirements are state or process requirements, I will not examine (iii) further.

When we turn to the fourth case, we find that the instrumental benefits of (iv) explainability in the individual condition are general as well as substantive, and that explanation via compliance with state requirements does not suffer from similar conceptual worries as prediction based on such requirements. As the benefits of explainability in the individual condition constitute the ultimate fundament of my suggested reason to be rational, supporting this tripartite claim is what will occupy us for the rest of this section.

Let me start with the third point. Because explanation is not about telling us which unknown attitudes we (will) also hold, but instead about telling us why we hold (or held) the known attitudes we hold (or held), the nature of state requirements does not block meaningful explanation in terms

¹⁴ If I am instead trying to predict what attitudes I will contemporaneously hold in the future, the worry largely stays the same. Cases in which I know that I will believe that p and that if p then q , but do not know whether I will also believe that q , although it matters to me whether q , seem marginal at best.

of such requirements. The explanation that I believe that q because I am rational and I also believe that p and that if p then q is meaningful because of the relation it highlights between my known attitudes. To talk of agents explaining their attitudes on the basis of state requirements is conceptually appropriate.

As for the second point, the *substantiality* of the benefits of being explainable, I take it that even if we are on our own having explainable attitudes is beneficial. Generally speaking, explainability is valuable because if something is explainable we can make sense of that thing. If something is unexplainable, this means that we can never answer the question ‘But why is it that...?’ with respect to that thing. Yet, if we cannot answer questions of this kind this means that we remain in a forced state of wonder.

Think of a toddler who has learned how to switch on a standard-lamp pushing the button on the back. He switches it on and off for some time, until he loses interest – and goes on to play with the lamp’s plug until, at some point, he pulls it out of the socket. If he then turns back to the switch in order to light the lamp, he will realize that nothing is happening. He will not understand why that is. Instead, he will look at the switch, work it, and wonder. If in this situation the parents do not intervene to explain to the toddler that the lamp can only be switched on if the plug is in the socket, as it was before, he is forced to remain in this state of wonder. The situation of the toddler is very much the situation we find ourselves in, if we encounter something unexplainable.¹⁵

Thus, being explainable in the individual condition is valuable as it allows us to leave a forced state of wonder and make sense of ourselves. And whereas understanding others in the social condition once again mostly seems to be of instrumental value, understanding ourselves intuitively is intrinsically valuable.¹⁶

¹⁵ It is true that sometimes we might actually want to remain in a state of wonder. For example, some people seem to think of romantic love as necessarily featuring an aspect of wonder as to why the beloved person returns one’s feelings. However, such voluntary states of wonder are very different from forced states of wonder.

¹⁶ In the context of this essay, it seems to be of particular importance that Kolodny thinks so as well (2005, 545–46).

Curiously, the question of why we should care about explaining anything at all often is not explicitly addressed even by those authors who talk about what an explanation is (Koertge 1992). Koertge (1992, 86–87) finds that, for instance, Hempel is surprisingly unclear about what we are trying to achieve with explanations. But Friedman, for one, gives a powerful answer in that he argues that the essence of explaining is to increase our understanding of the world (1974, 15). To explain things is to acquire comprehension of them, and comprehension is a valuable thing. According to Friedman, when we explain we render things intelligible in terms of other things by highlighting a unifying thread.

Explanation goes beyond answering what-questions. It does not merely describe things but – as every explanation is the answer to a why-question – makes sense of them. Explanation thus understood is sufficient for comprehension. Whenever we give a genuine explanation, i.e. an answer to a why-question which renders things intelligible in terms of other things, we have made sense of something. To be sure, there are better and worse explanations judging by different criteria. But the fact that explanations give answers that render things intelligible in terms of other things and thus provide comprehension is a characteristic shared by all explanations.

Tracing in the individual condition does not lead to comprehension *tout court*, but to comprehension of one’s own attitudes.¹⁷ But one might still wonder whether self-comprehension is intrinsically valuable, as I claimed it intuitively is. Consider the following example which might help to follow my intuition. Suppose I initially believe that q and that, for some reason, I reconsider whether I really believe that q . A natural question to ask myself in this situation seems to be: ‘Why did I believe that q again?’. That is, I start searching for an explanation of my initial belief. One

¹⁷ With respect to how we are able to make sense of ourselves if we are rational, compare Velleman’s (2000) argument that having transitive preferences promotes one’s ability to make sense of oneself. Because transitive preferences have a unifying thread the agent can see that they hang together and thus make more sense of himself than if they stood isolated in their multiplicity (2000, 158–60). In contrast to Velleman, the common thread I am suggesting is not constituted by compliance with the axioms of decision theory, but by compliance with rational requirements.

explanation I might come up with is the following. 'I believe that q because I'm rational and I also believe that p and that if p then q .' I have acquired self-comprehension via tracing.

It is important to realize what such self-comprehension does *not* give me. Just because I now understand why I believe that q , this does not tell me whether q . For instance, I might be subject to an error case. That is, self-comprehension does not lead to knowledge of the outside world. Still, I cannot help but feel that in this situation, understanding myself is a valuable thing. I now understand why I believe that q . That is something. It is precisely in light of self-comprehension falling short of knowledge of the outside world that we realize that self-comprehension already is valuable as such. I can be perfectly aware that self-comprehension does not have to lead to knowledge of the outside world and still feel that the self-comprehension I reached via tracing is valuable. This is because having an answer to my why-question concerning myself is valuable.

Let us finally proceed to the third point, the idea that the instrumental benefits associated to explainability in the individual condition defy concerns of *generality*. If we only look at the social condition, there indeed are instances where the agent does not have a reason to be traceable *for others*, like e.g. the diamond thief case. However, this does not automatically mean that the agent has no reason to be explainable *to herself*. First of all, notice that you can be untraceable to the outside while remaining traceable to yourself. This can be achieved, for instance, by hiding your beliefs and intentions from outsiders, thus not providing them with the initial conditions needed in addition to the general law in order to explain or predict. There is a distinction between our attitudes being untraceable for others because they lack knowledge of our attitudes, and our attitudes being untraceable for others because we are not coherent.

Furthermore, it seems especially important that at least the agent herself keeps track of her attitudes, when she is trying to hide her attitudes from outside observers. Of course, you could also be externally untraceable by being untraceable even to yourself. You would then be believing and intending in ways unexplainable or unpredictable for everybody. But this does not seem to be a particularly advisable strategy. A diamond thief unable to trace her attitudes most likely would

not be a very successful one. Put differently, in such cases the agent lacks reasons to be traced, but does not necessarily lack reason to be traceable.

We might also come up with scenarios in which we have a reason to be internally untraceable. If an evil demon threatens you to wreak havoc on the world unless you truly are unexplainable even to yourself – and if he has a way of telling – there is a substantive reason to be unexplainable. But still it is not quite the same to have a reason to be unexplainable and to have no reason to be explainable. Even in the evil demon scenario your reason for being explainable remains. Being explainable still is, *pro tanto*, a good thing. Its value is just counterbalanced, probably outweighed, by a reason for being unexplainable. Thus, the benefits of being explainable in the individual condition are general.

To summarize, I have argued that the benefits of traceability which underlie my reason to be rational seem to be most substantive in the social condition, but more general in the individual condition. The latter notably is the case because explanation via tracing in the individual condition is a sufficient condition for self-comprehension, which is intrinsically valuable. This is the ultimate fundament of my suggested instrumental reason to be rational as it ensures its generality. But importantly, we are not looking at two different reasons in the social and individual conditions. The reason to be rational, i.e. that being rational entails being traceable, remains the same in both conditions. Thus, the benefits we have from being traceable in the social condition add to those we have in the individual condition and make the suggested reason more substantive.

5 Objections

There are different worries that one might have with respect to how well my suggested reason meets the desiderata for a genuine reason to be rational as presented in Section 2. I will first address one objection to the generality of my suggested reason, then two objections to its substantiality. As my suggested reason to be rational is a direct reason to be rational, it does not rely on a potentially faulty principle of inference or some psychological impossibility (as e.g. ‘being fully

rational’). Thus, I do not have to further defend my suggested reason with respect to how it meets these latter two desiderata.

As an objection to the suggested reason’s generality, one might wonder how it fares with respect to error cases, i.e. with respect to the most prominent pitfall of many other candidate reasons. Addressing this worry, we find that it is a particularly noteworthy feature of the suggested reason to be rational that it is not foiled by error cases.

Consider the following example. If we know that an agent believes that p , and if we also know that he believes that if p then q , then we can explain his belief that q in terms of these two other attitudes. What happens if the agent believes that p although $\neg p$ (but is otherwise coherent)? In this case we can still explain without any problems that if the agent also believes that if p then q , he believes that q . Put differently, just because the agent believes p although $\neg p$ is true, this does in no way impede our ability to explain the agent’s belief that q as long as he is coherent.

Coherent reasoning is not paralyzed once we start out with a false belief – and neither is explanation. Our ability to explain an agent’s beliefs does not suffer from error cases; and neither does the agent’s ability to make sense of his own beliefs. The agent remains explainable to us and to himself, which is valuable. Thus, the suggested instrumental reason to be rational persists in error cases. This is because the agent’s explainability is based on the formal relation of his attitudes and entirely independent of their contents.¹⁸ For the exact same reasons as regards the explanation of an agent’s attitudes, error cases do not foil the ability to predict specific attitudes of a rational agent.

Here is a second objection to my suggested reason to be rational. It does not target the suggested reason’s generality, but its substantiality. Recall that tracing is a specific way of explaining

¹⁸ There might, of course, be reasons to believe $\neg p$ based on the content of $\neg p$. And sometimes these reasons might outweigh the benefits conferred by an agent’s attitudes complying with a certain formal structure. ‘Better to be incomprehensible, to myself and others,’ one might think, ‘than to believe *this*’. But this does not mean that there are no form-based *pro tanto* reasons to hold certain attitudes to begin with.

and/or predicting an agent's attitudes.¹⁹ If the reason to be rational is that being rational entails being explainable, then it is important to note that being rational is only one way that leads to explainability, but not the only way. Put differently, there are approaches other than complying with rational requirements which also establish explainability of our attitudes. Thus, it seems as if we might be able to secure the suggested benefit from being rational, namely being explainable, without being rational. Maybe we could even secure it in a better way.

What I am alluding to is that, in the social sciences, it is a well-documented fact that people *systematically* behave irrationally, e.g. when framing effects are involved (Tversky and Kahneman 1981). Whenever there are framing effects around, odds are that people will not behave strictly rationally. However, this does not impede our ability to explain their behavior. Precisely because these irrationalities are systematic in nature, we can come up with law-like generalizations in order to account for them. These law-like generalizations then assume the same role as the rational requirements do in tracing: the role of the general law we need in order to allow for deductive-nomological explanation.

There are hybrid theories which account for some systematic irrationalities and yet stress in which situations people typically do behave rationally. As these theories do not exclusively rely on the idea of coherence in order to explain or predict an agent's attitudes, but also integrate known exceptions, their explanatory and predictive power is supposedly greater than that of expecting pure coherence. Thus, if the suggested reason to be rational is that being rational entails having explainable attitudes, one might argue that we have a reason to comply with the regularities such hybrid theories of behavior identify rather than to comply with rational requirements. Put differently, in the light of alternative methods of achieving explainability, one might ask: what reason do we have to comply with rational requirements, if the reason to do so allegedly is to be

¹⁹ For the sake of brevity and in order to avoid unnecessarily cumbersome sentences, I will only use the terms 'explainability' and 'explainable' in order to refer to both explainability and predictability for the rest of this essay.

explainable, yet there are other ways to be, so to speak, even more explainable which do not rely on complying with rational requirements?

The short answer to this question is the following. My suggested reason to be rational, i.e. that being rational entails being traceable, is not affected by the existence of other ways of being explainable. A rational agent's attitudes are traceable whether or not there are other ways around to explain them. And traceability is associated to clear benefits. So we have a reason to be rational. Generally speaking, it is not the case that there cannot be two good explanations for the same thing, as one explanation might be better in some aspects whereas the other explanation might be better in other aspects or for a different purpose.

In contrast, the longer answer to the question addresses the worry that we might have *more* reason to comply with the regularities of a hybrid, 'mostly-rational-but-sometimes-systematically-irrational' theory of behavior than to straightforwardly comply with rational requirements, because the former seems to provide us with greater explanatory and predictive power. In my opinion, it is not quite clear whether we truly have more reason to do so, even if the ultimate reason to comply with either alternative (or any other theory able to explain an agent's attitudes) is explainability.

True, it seems that if we ultimately care about the explainability of an agent's attitudes, those approaches which give us a high amount of explanatory and predictive power *ceteris paribus* are preferable. From this point of view, it seems like one has more reason to let one's attitudes comply with the regularities identified by hybrid theories. Yet, the amount of explanatory and predictive power provided by a certain approach is not the only criterion to evaluate it. It is also true that, *ceteris paribus*, those approaches to making an agent's attitudes explainable which are simpler, more parsimonious, or more complete (in the sense of providing a clear explanation or prediction for every situation) are preferable. With respect to all three of these criteria, the approach of complying with rational requirements seems to score extraordinarily highly. With respect to parsimony, for instance, the theory of rationality makes far less assumptions than other alternatives. All it assumes is that an agent's attitudes comply with a finite set of coherence requirements. Put differently, even

if we ultimately care for the explainability of an agent's attitudes, the extent to which we reach this goal by a given approach should not exclusively be judged based on the criterion of explanatory and predictive power, but also based on other criteria such as simplicity, parsimony, completeness, or naturalness.

Let me be very clear about what I am not saying here. I am not arguing that being rational is the best way for an agent to be explainable. Nevertheless, being rational strikes me as a very good way for being explainable to oneself and to others. This is because rational requirements are a fairly simple, parsimonious, and complete set of rules to abide by. Furthermore, complying with rational requirements seems to be the *natural* or intuitive way of having explainable attitudes, for what that is worth. This latter claim is, of course, an empirical claim (and so are its refutations). But I think there are good ways to support it. For example, it seems widely acknowledged that accusing somebody of irrationality is a valid stand-alone criticism (cf. Kolodny 2005, 554). Whatever the content of our beliefs may be, if they do not coherently interrelate with each other, we will be criticized that even by our own lights we ought to believe something else – and that it would be unexplainable if we did not.

Again, the alleged naturalness of being explainable by being rational is not intended by itself to give a strong argument for why we have reason to comply with rational requirements rather than to adopt some other approach have explainable attitudes; at least not in the *individual* condition. In the *social* condition, however, the naturalness of being explainable by being rational is of some importance. Here is why.

In the social condition agents mutually try to explain and predict their attitudes in order to coordinate their behavior (Bratman 1987). In these situations, successful coordination is hugely facilitated if everybody shares the same idea of how to explain and predict an agent's attitudes. Just imagine two agents mutually trying to make sense of each other's attitudes with one agent supposing that the appropriate requirement attitudes should comply with is MP, whereas the other agent thinks the appropriate requirement is MP*.

*Modus ponens** (MP*): It is required of you that, if you believe p and you believe that if p then q , and if it matters to you whether q , then you believe $\neg q$.

In this case, in which the two agents refer to different requirements of how attitudes should be related amongst each other, they are not easily able to correctly explain or predict what the other will believe. Without further elaboration, they will not even notice why they disagree with respect to q , even if they both believe that p and that if p then q . Of course, if they figure out that they use different variants of *modus ponens*, then they are able to coordinate their attitudes. But in a world in which every agent follows a different approach to how attitudes should relate, people are not able to explain and predict each other's behavior efficiently.

Put differently, successful coordination of beliefs and actions in the social condition is way easier if there is a standard around for how we expect our attitudes and those of others to relate. For in this case no additional effort has to be exerted in order to identify the particular way of how to explain and predict a specific agent's attitudes. Whomever we interact with, we can simply expect that her attitudes will comply with the commonly accepted standard when explaining and predicting them. Certainly, it is even better if the commonly accepted standard is one that feels like the natural or intuitively correct one. For in this case comparably little effort has to be exerted in order to make people comply with it. If everybody thought that driving on the left side of the road is the natural or intuitively right approach to driving, no law would have to be passed and enforced to this effect. In the most extreme case, a certain standard can be so obvious and feel so natural to everybody, that it is self-enforcing.

However, in my opinion, compliance with the requirements of rationality just seems to be such a commonly accepted and intuitively appealing standard of how to explain and predict an agent's attitudes. The best way to argue for this claim seems to be to hold that rationality is a *focal point* (Schelling 1960) in this regard.

People can often concert their intentions and expectations with others if each knows that the other is trying to do the same. Most situations ... provide some clue for coordinating behavior, some *focal point* for each person's expectation of what the other expects him to expect to be expected to do. (Schelling 1960, 57–58)

A focal point is a standard *malgré lui*. Nobody established it – it just emerged spontaneously. The classic example of a focal point is the answer that New Yorkers give to the question where they would expect to meet up with somebody given they had not had the chance of agreeing on a place to meet in advance. Most of them reply: (at the info booth) at Grand Central Station (1960, 57–58). There just is something about Grand Central Station that makes it the unique natural or prominent place to meet. In much the same way I think that conformity to rational requirements is a focal point with respect to how to explain and predict an agent's attitudes.²⁰ There is something about rationality that makes it a natural or prominent approach for explanation.²¹ Maybe it is the

²⁰ Note that conceiving of rationality as a focal point is compatible with some experts in the behavioral sciences expecting people to behave systematically irrationally in specific situations. Schelling highlights that the prominence or naturalness of focal points 'depends on time and place and who the people are. Ordinary folk lost on a plane circular area may naturally go to the center to meet each other; but only one versed in mathematics would "naturally" expect to meet his partner at the center of gravity of an irregularly shaped area' (1960, 58). Thus, if we conceive of rationality as a focal point, experts may well expect that rationality sometimes is systematically violated, whereas the huge majority of 'ordinary folk' expects coherent attitudes throughout.

²¹ Notice that I am not suggesting here that humans have a *disposition to be rational* and that *this is why* there is a reason to be rational. I am decidedly not adopting the strategy Broome (2007a) discusses as described in Section 2; namely to first show that we ought to have a disposition to be rational that makes us satisfy many rational requirements, and then infer that, therefore, we have an indirect instrumental reason to be rational. I think that we indeed have a fairly stable disposition to be rational. But first, I do not say anything concerning whether we ought to have this disposition. And second, the disposition is not why there is (indirect) reason to be rational. Instead, the reason to be rational that I am suggesting is a *direct* instrumental reason: being rational entails being traceable. I am only alluding to the idea of

close relationship rational requirements have to fundamental principles in logic – which one may deem to be self-evident. Although it is theoretically possible for there to be a different commonly accepted standard other than abidance by rational requirements, e.g. one that integrates Tversky's and Kahneman's findings, this just seems to be empirically false. Therefore, and for all the other reasons I have given, I think that being rational is a very good way to be explainable – even if it might not be the best. My suggested reason to be rational, i.e. that being rational entails being traceable (and thus explainable), holds its ground in light of alternative explanatory approaches.

Here is the third and final objection. One might object not that there are better ways to be explainable other than compliance with rational requirements, but that there are cases in which such compliance does not necessarily allow us to explain an agent's attitudes at all. The idea is that, for instance, if we are explaining an agent's belief that q in terms of her beliefs that p , and that if p then q , and her compliance with MP, this might not tell us *why* she believes that q . This is because she could believe that q for some altogether different reason. And if compliance with rational requirements seemingly has no explanatory role for her belief whatsoever, it is hard to see how explainability via traceability would lead to valuable comprehension. Consider the following example to illustrate the objection.

RAIN: An agent believes that it has been raining (p), as she is sitting in front of a window of a room overlooking a park, and she has seen the rain falling for quite some time. She also believes that if it has been raining, the streets are wet (if p then q). A friend of hers enters the room and tells her that 'The streets are wet', pointing to his ruined new shoes.

a disposition to be rational in order to suggest an idea of why we might feel some inner and outer pressure to be rational and why we tend to regard rationality as an intuitive standard for explaining and predicting attitudes. But I am not arguing that our having this disposition directly is a reason to be rational.

Suppose the friend is an exceptionally trustworthy friend, and she always believes whatever he tells her. (Her trust in his statements is a general law.) In this scenario, one might argue that even if the agent believes that p , and that if p then q , and if she is complying with MP, this does not necessarily explain her belief that q – if what explains her belief that q is that her friend told her so. And one might add that if compliance with MP has no explanatory role in RAIN, then the suggested instrumental reason to be rational no longer exists.

This third objection seems more pressing if we conceive of rational requirements as state rather than process requirements. Let me first outline why that is, and how the suggested reason to be rational is not thwarted if rational requirements are process requirements, before turning to discussing state requirements.

Following Kolodny, '[p]rocess requirements require you to *do* something over time, where “do” is understood broadly, so as to include forming and revising beliefs' (2007, 371); they tell us 'how, going forward, one is to form, retain, or revise one's attitudes so as to avoid or escape ... conflict-states' (2005, 517). Thus, if we re-conceive MP as a process requirement, it does not merely tell us that rationality requires that if we believe that p , and that if p then q , and if it matters to us whether q , that we also believe that q . Instead, it tells us, for example, that rationality requires that – if we believe that p , and that if p then q , and if it matters to us whether q – we either (i) form the belief that q (in case we had not made up our mind about q before) or *revise* our belief that $\neg q$, or (ii) *drop* the belief that p , or (iii) *drop* the belief that if p then q . Basically, whereas state requirements tell us what state our attitudes ought to be in for us to be rational, process requirements tell us what to do whenever our attitudes are not in a rational state as defined by state requirements.

Because state and process requirements differ in this way, it means something different to comply with state or process requirements, respectively. As state requirements merely require that your attitudes be in a certain state, compliance with a state requirement is not undermined if you reach this state 'in the wrong way'. On the other hand, as process requirements tell you what you

are required to do over time, they are less open in this respect. Arguably, you do not comply with a process variant of MP if, say, you believe that p , and that if p then q , and if it matters to you whether q , yet you only form the belief that q because your friend tells you that q and thus makes you believe that q . In this scenario, you arguably do not comply with a process variant of MP, as your new belief that q has not been formed ‘in the right way’ in order to comply with that requirement. Notably, doing what rationality requires has not been the ‘primary reason’ (Davidson 1963, 686) for forming the belief that q .

Interestingly, although Broome (2007b, 366–68) is skeptical about rational requirements being process requirements, his criticism of Kolodny’s (2007) argument in favor of process requirements reveals how both Broome and Kolodny understand process requirements in the way just outlined.²² Both seem to think that compliance requires actually following a process that leads us ‘in the right way’ from antecedent-beliefs to consequent-belief, so that you form the consequent-belief because of your compliance. Broome, for instance, highlights that the process required by process requirements cannot simply be ‘to do some reasoning based on the [antecedent-beliefs]’ (2007b, 367). For if it were, then a process variant of MP would not strictly speaking *require* that, if you believe that p , and that if p then q , and if it matters to you whether q , that you also form the belief that q . In his charitable effort to reconstruct Kolodny’s argument, Broome reaches the conclusion that the process in question must be a ‘process such that, necessarily, if you ... [follow the process] you [form the belief that q] at a time later than t ’ (2007b, 368) – with t being the time at which you hold the antecedent-beliefs that p , and that if p then q .

Thus, if we conceive of rational requirements as process requirements in the same way Broome and Kolodny do, we can distinguish two scenarios with respect to RAIN. Either the agent complies with a process variant of MP. If so, then her compliance with MP also tells us something about *why* she believes that q ; because compliance means to form the belief that q because of following

²² Way (2011, 238) suggests a more lenient conception might be appropriate in some circumstances.

the process set out in the process variant of MP. So the objection that she instead believes that q because of her friend's testimony is a non-starter. Or the agent forms the belief that q because of her friend's testimony; in which case she no longer complies with a process variant of MP. Likewise, the objection does not get off the ground. For it is no longer the case that she complies with a process variant of MP yet something else explains her belief that q .²³

If we conceive of rational requirements as process requirements, we cannot really raise an objection like the third one. However, if we conceive of rational requirements as state requirements, objection three seems forceful; at least *if we target the individual condition*.

Note that *in the social condition* it is widely accepted in many social sciences (and notably in economics), that explanations of other people's attitudes and actions do not even have to be factive, or correspond to the explanation the agent gives herself if asked, in order be instrumentally valuable. To cite Friedman's (1953, 20) famous example: even if a snooker player does not actually calculate the angles of his next shot and the required spin of the ball in his head, we can explain and predict his behavior in terms of the geometry and physics of snooker. Whenever we rely on models to explain and predict, we do not tend to think that we are giving an 'actual' or factive explanation of an agent's attitudes. Still, such explanations are highly useful for explanatory and

²³ It is noteworthy that neither Broome (2007b) nor Kolodny (2005, 2007) argues in favor of rational requirements being wide-scope process requirements. Broome is skeptical of process requirements of rationality in general, and Kolodny criticizes *wide-scope* process requirements in particular. Yet, this is not the right place to settle the controversy over whether rational requirements are (wide-scope) process or state requirements. Suffice it to say that, first, Kolodny is certainly right if he says that 'our ordinary attributions of irrationality are at least sometimes about what people do, or refuse to do, over time' (2007, 371). So there is a natural language argument in favor of there being *process* requirements of rationality – just as there is, according to Broome (2007b, 361–62), a natural language argument in favor of rational requirements being *wide-scope*. Second, neither Broome's skepticism about process requirements, nor Kolodny's criticism of wide-scope process requirements are entirely conclusive in the sense that they would strictly rule out the existence of wide-scope process requirements of rationality. For an argument to counter Broome's skepticism see Kolodny (2007, 381–82). For an argument against Kolodny's criticism see Way (2011, 237–38).

predictive purposes. Thus, in the social condition, the instrumental value of being traceable is not affected by the objection that the agent believes that q for a different reason.²⁴

So the objection only stands strong in the individual condition, and if we conceive of rational requirements as state requirements. If the agent in RAIN tries to explain her belief that q , aiming for self-comprehension, we might feel that compliance with state requirements fulfills no explanatory role and is thus unable to ground a reason to be rational.

Now, in order to reply to this, it is crucial to realize that RAIN actually is a case of *overdetermination*. In RAIN, there are two deductive-nomological explanations which are both sufficient to explain the agent's belief that q . (i) If she always believes what her friend tells her, and if he tells her that q , this sufficiently explains her belief that q . (ii) If she is complying with MP, and if she believes that p , and that if p then q , this also sufficiently explains why she believes that q . Since both explanations are factive and sufficient, we can hardly keep up the idea that compliance with rational requirements plays *no* explanatory role for the agent's belief that q in RAIN. So if, as stated above, the objection is that compliance with rational requirements does not necessarily allow us to explain the agent's belief that q , that is actually wrong. Rather, we are dealing with a case in which we necessarily have a sufficient explanation for the agent's belief if the agent complies with MP – only that we do not necessarily need it. For if the agent explains her belief that q in terms of her friend's statement, there is not strictly speaking a need for an additional sufficient explanation. The agent has already achieved self-comprehension.

I can understand why in RAIN one might want to say that the agent does not believe that q because of complying with MP. For even if it matters to her whether q , it surely feels like q has come to the forefront of her attention because of her friend's testimony. Phenomenologically

²⁴ Dancy gives another example of a non-factive explanation in the social condition, which, as he argues, yields valuable comprehension: 'His reason for doing it was that it would increase his pension, but in fact he was quite wrong about that' (2000, 132).

speaking and from her own perspective, although she is complying with MP all along, it does not seem to be her compliance with MP that explains why she believes that q . But the fact of overdetermination remains. To be sure, whenever there is more than one sufficient explanation, we seem to have an urge to identify the single ‘right’ explanation. But as it is vain to try to identify the single ‘right’ explanation for why the victim dies in firing squad cases, so it is in RAIN.

We do not strictly speaking *need two* sufficient explanations in cases of overdetermination, if we already have one that establishes self-comprehension. But this does not mean that a second sufficient explanation is not valuable. This is because – unlike in legal contexts in which we might struggle to ascribe individual responsibility to the members of a firing squad – overdetermination in explanatory contexts is not a nuisance. It is something we should welcome when we are aiming for comprehension.

As briefly touched on above, the essence of explanation is to increase comprehension. Comprehension, in turn, arises if we render things intelligible in terms of other things by highlighting a unifying thread (Michael Friedman 1974; Kitcher 1981).²⁵ Based on this unificationist conception of comprehension, overdetermined explananda are not problematic, but might be understood as instances of *deeper* comprehension. If comprehension is to see how things hang together, then encountering that they hang together in multiple ways is desirable; a healthy sign of deep comprehension. If somebody makes us aware that – by our own lights! – we do not only have one but two sufficient reasons for why we believe something, this surely feels like our self-comprehension is deepening. It might be that because the debate over the normativity of rationality typically focuses on *local* (ir)rationalities we too readily put aside the fact that we have entire networks of beliefs and intentions, which hang together in multiple ways. Overall, it seems

²⁵ This unificationist account of comprehension developed out of the roots of the Hempel and Oppenheim model of deductive-nomological explanation .

that with respect to an agent's attitudes, all kinds of overdeterminations are the rule rather than the exception.

To be sure, if we find ourselves in the individual condition in a case like RAIN, then the value of being explainable via being traceable is small. One might have hoped for a more substantive reason to be rational than deepened self-comprehension. Still, my reason to be rational withstands this final objection even if rational requirements are state requirements.

References

- Bratman, Michael. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- . 2009. "Intention, Practical Rationality, and Self-Governance." *Ethics* 119 (3): 411–43.
- Broome, John. 1999. "Normative Requirements." *Ratio* 12: 398–419.
- . 2007a. "Is Rationality Normative?" *Disputatio* 2 (23): 161–78.
- . 2007b. "Wide or Narrow Scope?" *Mind* 116: 359–70.
- . 2013. *Rationality through Reasoning*. Chichester and Malden, MA: Wiley Blackwell.
- Brunero, John. 2010. "The Scope of Rational Requirements." *Philosophical Quarterly* 60: 28–49.
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Davidson, Donald. 1963. "Actions, Reasons, and Causes." *The Journal of Philosophy* 60 (23): 685–700.
- . 1973. "Radical Interpretation." *Dialectica* 27: 314–28.
- Friedman, Michael. 1974. "Explanation and Scientific Understanding." *The Journal of Philosophy* 71 (1): 5–19.
- Friedman, Milton. 1953. "The Methodology of Positive Economics." In *Essays in Positive Economics*, 3–43. Chicago: University of Chicago Press.
- Kitcher, Philip. 1981. "Explanatory Unification." *Philosophy of Science* 48 (4): 507–31.
- Koertge, Noretta. 1992. "Explanation and Its Problems." *British Journal for the Philosophy of Science* 43 (1): 85–98.
- Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114 (455): 509–63.
- . 2007. "State or Process Requirements?" *Mind* 116 (462): 371–85.
- Reisner, Andrew E. 2009. "Unifying the Requirements of Rationality." *Philosophical Explorations* 12 (3): 243–60.
- . 2011. "Is There Reason to Be Theoretically Rational." In *Reasons for Belief*, edited by Andrew E. Reisner and Asbjørn Steglich-Petersen, 34–53. Cambridge: Cambridge University Press.
- Schelling, Thomas C. 1960. *The Strategy of Conflict*. London: Oxford University Press.

- Schroeder, Mark. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18: 337–64.
- Tversky, Amos, and Daniel Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice." *Science* 211 (4481): 453–58.
- Velleman, J. David. 2000. "The Story of Rational Action." In *The Possibility of Practical Reason*, 144–69. Oxford: Clarendon Press.
- Way, Jonathan. 2011. "The Symmetry of Rational Requirements." *Philosophical Studies* 155 (2): 227–39.