



Requirements of intention in light of belief

Carlos Núñez¹ 

© The Author(s) 2019

Abstract Much work in the philosophy of action in the last few decades has focused on the elucidation and justification of a series of purported norms of practical rationality that concern the presence or absence of intention in light of belief, and that demand a kind of structural coherence in the psychology of an agent. Examples of such norms (all roughly formulated) include: Intention Detachment, which proscribes intending to do something in case some condition obtains, believing that such condition obtains, and not intending to do that thing; Intention-Belief Consistency, which proscribes intending to do what you believe you will not do; Intention Consistency, which proscribes intending each of two ends you believe to be inconsistent; and Means-End Coherence, which proscribes intending an end and not intending the means you believe to be implied by your end. In this paper, I present a series of examples that show that these requirements are not genuine requirements of rationality. The reason for this is simple: these requirements concern the presence or absence of intention in light of all-out belief. Rational agents like us, however, do not, and in fact should not, always form or revise their intentions in light of what they all-out believe. When such agents do not form or revise their intentions in light of what they all-out believe, they need not be irrational if they do not conform to these requirements.

Keywords Practical rationality · Intention · Belief · Intention consistency · Means-end coherence · Enkrasia · Coherentism · Myth-theory

✉ Carlos Núñez
carlos.nunez@univie.ac.at

¹ Department of Philosophy, University of Vienna, Ebendorferstraße 10/13, 1010 Vienna, Austria

1 Introduction

Much work in the philosophy of action in the last few decades has focused on the elucidation and justification of a series of purported norms of practical rationality that concern the presence or absence of intention in light of belief, and that demand a kind of structural coherence in the psychology of an agent. Some examples, roughly formulated, include the following.¹ It is commonly thought that rationality requires of an agent A that:

- *Intention Detachment*: A does not [intend that (*p if q*), believe that *q*, believe that (were she not to intend that *p*, then because of that *not p*), and not intend that *p*].
- *Intention-Belief Consistency*: A does not [intend that *p* and believe that *not p*].
- *Intention Consistency*: A does not [intend that *p*, intend that *q*, and believe that (*p* and *q* are inconsistent—in the sense that, were it to be the case that *p*, then because of that *not q*, or vice versa)].²
- *Means-End Coherence*: A does not [intend that *p*, believe that (intending that *q* is a means implied by *p*—in the sense that, were she not to intend that *q*, then because of that *not p*), and not intend that *q*].³

Interest in these norms has been fueled in part by a move away from a simple belief-desire model of human agency and its focus on the maximization of expected utility as the sole principle of practical rationality, towards a more complex model that incorporates intentions as genuine attitudes, characterized by their own distinctive functional features and rational demands.

Let me call purported norms of practical rationality that demand structural coherence as such, independently of any other concern, “requirements of coherence,” and proponents of them “coherentists” about practical rationality. Moreover, let me call the specific set of requirements just mentioned requirements of “Intention-Belief coherence” (IB-coherence), and proponents of them “Intention-Belief coherentists” (IB-coherentists).⁴

The coherentist project has not gone unchallenged. Different theorists have argued that it is a mistake to think that rationality requires structural coherence as such. They have different reasons for thinking this. For example, some people

¹ I say “roughly formulated” in part because all these norms need appropriate references to times. I shall not make such references explicit here. The requirements are ‘wide-scope’ in the sense made familiar by Broome (e.g. 2013). For clarity, I mark the scope of the requirement with brackets. Similarly, I put the content of beliefs or intentions in parenthesis when it is complex.

² Throughout, I use the term ‘inconsistency’ exclusively in this sense.

³ I take the term ‘means implied’ from Broome (2013).

⁴ Of course, both IB-coherentists, and coherentists more generally, need not deny that there are other requirements of rationality that go beyond requirements of formal coherence. Two notorious IB-coherentists, in the specified sense, are Bratman (1987) and Broome (2013).

believe that practical rationality consists solely of maximizing expected utility.⁵ Others think that it consists solely of responding correctly to the reasons you have, or, alternatively, to the reasons you believe yourself to have, or that are somehow made available by your evidence.⁶ Because of this, they believe that whether or not an agent is coherent is, at best, only indirectly related to the question whether she is rational. What really matters, rationally speaking, is maximizing expected utility, or responding correctly to (believed/available) reasons. I will call proponents of this idea “myth theorists” about coherence requirements of practical rationality.

Here, I will present a series of examples that show that, indeed, the requirements mentioned above are not genuine requirements of rationality.

The reason is simple: the listed requirements concern the presence or absence of intention in light of all-out belief. Rational agents like us, however, do not, and in fact should not, always form or revise (or, as I will also put it, ‘regulate’) their intentions in light of what they all-out believe. When such agents do not regulate their intentions on the basis of what they all-out believe, then breach of these requirements need not imply irrationality.

Here, I assume that a requirement is a genuine requirement of rationality only if, necessarily, breach of it implies irrationality.⁷ That is, I assume that genuine requirements of rationality impose a kind of “strict liability” with respect to rationality. Following Broome (2013), let me call this property “rational strict liability”:

- Rational strict liability: requirement R imposes rational strict liability iff, necessarily, if R requires of A that p , and it is not the case that p , then A is irrational.

In what follows, I will present a series of examples where agents are in breach of each of the listed requirements and yet are, intuitively, not irrational. Under the assumption that genuine requirements of rationality impose rational strict liability, this would show that these are not genuine requirements of rationality.⁸

This would make it seem like I am arguing in favor of a myth theory. But, in fact, I think coherentists should embrace this conclusion. This doesn’t mean coherentists should become myth theorists. It simply means that they need a more nuanced

⁵ Some of the things McCann says in (1986, 1991) suggest he believes something along these lines. Since he does not formulate his view exactly in these terms, however, I hesitate to attribute this position to him.

⁶ See e.g. Raz (2005), Kolodny (2005, 2007, 2008a, b) and Kiesewetter (2017).

⁷ When I speak of ‘irrationality’ I simply mean lack of perfect rationality. An agent who is irrational, in this sense, is just not perfectly rational.

⁸ The assumption that genuine requirements of rationality impose rational strict liability is not uncontroversial. It is endorsed by Broome (2013), but it might be legitimately questioned. Take Intention-Belief Consistency as an example. As David Hills pointed out to me (in conversation), one might think that one can rationally retain both the belief that *not p* and the intention that *p* in circumstances in which nothing yet favors giving up one over the other, and where it would be defeatist then to give up both, and still maintain that one is *under rational pressure* to resolve this conflict by finding appropriate grounds to give up one or the other. If my central suggestion in this paper works, however, it would also work against a view that understands rational requirements as imposing this weaker condition, since it would show that there might be cases where the agent holds the relevant attitudes and there isn’t even this rational pressure to find grounds to give one of them up.

picture of what practical rationality requires of an agent when her intentions are not regulated in light of her all-out beliefs.

I will proceed as follows. In the next section, I will mention the assumptions about the nature of belief on which I will be relying. In Sect. 3, I present counterexamples to each of the aforementioned requirements. In Sect. 4, I explain what lessons I think we should draw from these cases. In Sect. 5, I consider the special case of another requirement of coherence that also concerns the presence of intention in light of belief, but that seems to be surprisingly immune to the problem that plagues all the other requirements considered. This is Enkrasia, which demands (roughly) that you intend to do what you believe you ought to do. I will explain why I think Enkrasia is in this sense immune.

2 Belief and certainty

The cases I will present depend on the following relatively uncontroversial ideas, which I will assume for the purposes of this paper. The first idea is that it doesn't make rational sense to be willing to bet on the truth of a proposition p , at any odds, if one is not certain that p .⁹ This is because, if one is uncertain that p , then there will be bets where the expected utility of taking them would be outweighed by the expected utility of rejecting them. Suppose I am .9 confident that p . Then it would be irrational for me to take any bet with odds worse than 1:9 on the truth of p . I take this idea to be uncontroversial.

The second idea is that belief—by which I mean “all-out belief”—does not imply certainty. That is, one can believe that p without being certain that p . This idea is not completely uncontroversial.¹⁰ But it does represent the consensus and is, for familiar reasons, quite burdensome to reject.¹¹ Doing so implies denying that people believe most of the things we take them to believe, since people are not certain of most of the things we take them to believe. It also implies denying that people rationally believe most of what we take them to rationally believe, since we think that epistemic justification for certainty that p would require evidence that rules out any possibility that *not* p , whereas we allow that one could be epistemically justified in believing that p even when one's evidence does not completely rule out the possibility that *not* p .

I think these reasons are enough to reject the idea that belief implies certainty. If, however, one accepts certain familiar Bayesian principles, then there would be two further (though obviously related) reasons to reject this view. First, the view would imply that one should be willing to accept any bet on the truth of one's beliefs, regardless of the odds. After all, to be certain that p is to disregard any possibility that *not* p and, consequently, any undesirable outcome on the contingency that *not* p . So, if one believes that p , one should be willing to stake everything one has,

⁹ By “certain” I mean having a credence of 1.

¹⁰ For example deFinetti (1970), Levi (1991), Clarke (2013) and Dodd (2017).

¹¹ The idea that belief does not imply certainty is the general consensus and so is normally simply assumed. But it has been defended, to a greater or lesser extent, by some theorists. Some examples include Maher (1986, 1993), Kaplan (1998), Christensen (2007), Buchak (2013), Frankish (2004, 2009), Fantl and McGrath (2002, 2009), Worsnip (2016) and Hawthorne et al. (2016).

however valuable, against any possible payoff, however insignificant, on the truth of p . Second, it would imply that it would never be rationally permissible to revise one's beliefs: that, once one comes to believe that p , there is no piece of evidence one could ever come to be aware of, no matter the kind or amount, that could rationally allow one to change one's mind on that issue. This is because, given a credence of 1 that p , there is no way to update by conditionalization to a credence of less than 1 on that proposition.¹² For these reasons, I assume from now on that belief does not imply certainty.¹³

Together, these two ideas strongly suggest that there can be cases where it would make no rational sense to deliberate, and so to regulate one's intentions, on the basis of one's beliefs.¹⁴ This is because deliberating on the basis of a belief (by which I mean: evaluating the relevant alternatives conditional on the truth of the belief) is, in a natural sense, a way of betting on its truth. But if belief does not imply certainty, and sufficiently unfavorable odds can make it irrational to bet on the truth of a proposition if one is not certain of its truth, then this suggests that there can be cases where it would be irrational to deliberate—and so to regulate one's intentions—on the basis of what one believes.¹⁵

We are familiar with such cases. Oftentimes what we do and rationally ought to do is deliberate on the basis of the probabilities we assign to the relevant contingencies. Other times, what we do and ought to do is deliberate on the basis of

¹² Although see Weisberg (2009).

¹³ Clarke (2013) thinks he can avoid these problems by construing belief as credence 1 *in a context*. I find this implausible. Intuitively, an agent can believe that p while recognizing, within the same context, that there is a non-zero chance that *not* p (e.g. I believe I have hands, but I recognize, in the same context, that there is a non-zero—though normally negligible—chance that I do not). On this and other reasons against Clarke's proposal, see Worsnip (2016).

¹⁴ On this point, see especially Maher (1986, 1993) and Kaplan (1998).

¹⁵ Given our cognitive limitations, it is plausible to think that all-out beliefs will always, inevitably, play a role in somehow framing the space of possibilities that we consider in practical deliberation. We simply do not have the computational capacities to do away with all-out beliefs. Deliberative questions are typically framed against the background of propositions that are taken as certain for the purposes of solving the question at issue. This is so even in those cases where we are explicitly calculating the expected utility of possible actions based on our estimation of the probabilities of different contingencies. Among other things, the expected utility of the available acts is calculated by way of treating the consequences of such acts as certain in a given contingency. For example: if I am wondering whether to carry an umbrella or not, I may deliberate on the basis of the probabilities I assign to it raining or not, but I will typically treat the consequences of my acts as certain in a given contingency. That is, I will typically take it as given, for example, that if I carry an umbrella and it rains, I will not get wet; or that, if I don't carry an umbrella and it rains, I will get wet. I all-out believe each of these propositions, and because of this I frame the problem I face on their basis. However, I do not have a credence of 1 on them. I recognize that there is a non-zero chance of getting wet when it rains even though I do carry an umbrella, or of not getting wet when it rains even though I do not carry an umbrella. Still, even if I am not certain of them, I do deliberate on their basis. If instead I tried to deliberate exclusively on my assignment of probabilities, the problem would become very quickly unmanageable. This point is recognized by decision theorists. As Savage puts it, "we must expect acts with actually uncertain consequences to play the role of sure consequences in typical isolated decision situations" (1972, p. 84). (See also Joyce 1999; Ross and Schroeder 2014). Although *cf.* Maher (1986)). Still, even if we can't do away with all-out belief altogether, it is true that sometimes we do not, and rationally should not, deliberate on the basis of *certain* beliefs, even when those beliefs are relevant to the deliberative question at issue.

propositions we, for whatever reason, accept in a certain context, regardless of whether we believe them or not.¹⁶

Consider: I believe I am blood-type O-positive. I am not completely certain, but I clearly remember different doctors telling me so. Now contrast two deliberating contexts in which I am asked about my blood type, where I have the option of either simply answering that I am O-positive based on my belief, or of taking a blood test, offered by the person asking the question, that would provide the sure answer but cost me \$10. In one context, the question is asked by a student for the purposes of gathering statistical information. In the other, the question is asked by a paramedic who is about to give me a blood transfusion. Since nothing too bad would happen if I say to the student that I am O-positive in the contingency that I am not, whereas something terribly bad would happen if I say to the paramedic that I am O-positive in case I am not, the expected utility of deliberating on the basis of my belief is greater than that of deliberating on the basis of the relevant probabilities in the survey scenario, but lower in the transfusion scenario. If this is so, then, plausibly, in the one case I would and should deliberate on the basis of my belief and simply say that I am O-positive without taking the test, whereas in the other I would and should deliberate on the basis of the probabilities I assign to the relevant contingencies and take the test instead.

Obviously, conditional on my being O-positive, answering without taking the test is in both scenarios the dominant strategy, because it costs me \$10 to take the test and nothing not to take it. But it would be foolish of me, in the transfusion case, to evaluate the relevant alternatives conditional on the proposition that I am O-positive, since I am not sure that I am, and that so much is at stake. This doesn't show that I don't *really* believe I am O-positive. It just shows that there are cases where what is at stake might make it too risky for me to deliberate on the basis of propositions of which I am not sure (or, in any case, not sure enough).

Naturally, the question of whether my unwillingness in such a scenario to deliberate on the basis of that proposition shows that I do not really believe it, depends on what we say all-out belief is, and there are different possibilities here. As far as I can tell, though, there is no plausible theory of belief according to which such unwillingness implies lack of belief. Since the issue is important for my purposes, though, let me briefly consider some of those possibilities:

One possibility is that all-out belief just *is* certainty.¹⁷ Since I am assuming that belief does not even imply certainty, I shall ignore this view.

Another possibility is that belief is a contextually invariable degree of confidence. Suppose we say the degree is anything above .9. Then we could stipulate that I am .91 confident that I am O-positive. Still, given the odds, I wouldn't and shouldn't deliberate under the assumption that I am, even though, by stipulation, I do believe I am.

Some people have suggested instead that belief is a contextually variable degree of confidence, where the degree corresponds, precisely, to how confident you would

¹⁶ On this last point, see especially Bratman (1992) and Alonso (2014, 2016).

¹⁷ For example, Clarke (2013) and Dodd (2017).

need to be of a proposition to be willing to deliberate, in a given context, under the assumption of its truth.¹⁸ By stipulation, then, in the transfusion case I do not believe that I am O-positive, since I am not confident enough to deliberate under this assumption.

Now, it is not my aim here to argue against this view in any detail, but let me just briefly say why I think we should reject it. The most basic reason is that it conflicts with a plausible and firmly established view of the nature of belief and its overall role in rational agency. One way to formulate the issue is to point to the familiar idea that belief somehow ‘aims’ at truth.¹⁹ Needless to say, it is difficult to specify precisely what this amounts to, and I won’t try to do so here, but it involves ideas roughly along the lines that belief is, at the functional level, ‘regulated’—so ‘formed, revised and extinguished’ to use Velleman’s (2000) phrase—in response to evidence of truth, and that it is, at the normative level, assessable in relation to truth, both objectively (so that a belief is correct iff it is true) and subjectively (so that a belief is rational or warranted to the extent that it is appropriately regulated in accordance to standards that would be, somehow, conducive to truth). The view under consideration conflicts with this familiar picture because, under it, belief comes to be regulated and assessable by truth-independent factors, like what the odds are in a given deliberative situation, or even by such outlandish factors as whether you have had breakfast or not, since such factors might affect your willingness to take risks.

This is a radical departure from our ordinary concept of belief, and it has some rather unpalatable implications. Just to illustrate: I presume you believe you have a head. I also presume that you are extremely confident of this. Under the present view, I could deprive you of this belief by offering you a bet at sufficiently unfavorable odds on this proposition. And this regardless of the fact that you would (we can assume) remain as confident and as thoroughly convinced as ever that you do, in fact, have a head, and that your evidence for this has not changed one bit and is as conclusive as ever. I presume you also believe you have hands. If so, given the deliberative questions you face in the situation in which I offer the bet, you believe you have hands but you do not believe you have a head, even though (we can assume) you are equally confident that you have a head as that you have hands; and even though you recognize that the evidence for, and the probability of, the two propositions is exactly the same; and even though—and this point is crucial—the attitudes you have towards these two propositions play *exactly* the same kind of functional role in your overall psychology (e.g., you are disposed to regulate them in response to the same kind of evidence, to deliberate on their basis relative to exactly the same odds, stakes, deliberative questions, etc.). I find this incredible. Of course, implications like these do not *prove* that this view is false. But they do show it to be rather costly, and far enough removed from our ordinary concept of a belief that I

¹⁸ The view I am now considering is a view about the nature of belief, not one about the nature, exclusively, of *rational* belief. I interpret Weatherston (2005) and Ganson (2008) as presenting views of this kind. I think it is also possible to interpret Fantl and McGrath (2002, 2009) in this manner, although perhaps they should be interpreted instead as presenting a view of rational belief only.

¹⁹ See e.g. Velleman (2000), Wedgwood (2002), Shah (2003) and Shah and Velleman (2005).

think we would do well to avoid it unless the reasons for it are compelling.²⁰ As far as I can tell, however, there aren't very good reasons to adopt it.²¹ Because of this, I will from now on ignore it.

Another possibility, or family of possibilities, is to say that all-out belief is a kind of intentional disposition. Here are some options:

One is that a belief that p is a non-defeasible disposition to deliberate on the basis of p . Since this is a non-defeasible disposition, it entails that all-out belief implies certainty, and so can be ignored for the same reasons as before.

One might say instead that belief that p is a *defeasible* disposition to deliberate on the basis of p . This view avoids the previous problem. However, under most plausible ways of specifying what the defeating factors might be, the stakes or odds of a given deliberative context will most certainly figure among them (as is actually the case, among the views of this kind on offer).²² But if this is so, then, once again, I might be defeasibly disposed to deliberate under the assumption that I am O-positive, and yet be unwilling to deliberate on the basis of this belief in the transfusion case, because my disposition is defeated by the odds in that context.

Another option is to say that all-out belief is a linguistic disposition to assert the proposition believed. Mark Kaplan gives us a specific version of this idea. According to him: "You count as believing P just if, were your sole aim to assert the truth (as it pertains to P), and your only options were to assert that P, assert that \sim P or make neither assertion, you would prefer to assert that P." (Kaplan 1996, p.109) Well, we could stipulate that this is true in the transfusion case. If my sole aim were to assert the truth, and my only options were to assert that I am O-positive, that I am not, or to say nothing at all, I would prefer to assert that I am O-positive. Still, I might rationally refuse to rely on this belief and choose to take the blood test instead.

No doubt there are other views concerning the nature of all-out belief. But I know of no plausible view that implies that my not being willing to answer to the paramedic without taking the test implies that I do not believe I am O-positive. So I will assume from now on that it does not.

²⁰ Again, these considerations target a view about the nature of belief, not rational belief. As an anonymous referee points out, one could maintain that only truth related considerations provide reasons for belief, while maintaining that pragmatic considerations affect the amount of evidence needed for *warranted* or *justified* belief in a given context. According to a view of this kind, the fact that you are not willing to deliberate in a given context on the proposition that p doesn't imply that you don't believe in that context that p . Because it doesn't imply this, it is compatible with the central claim of this paper. However, it does affect some of the things I go on to say. I explain how in footnote 24.

²¹ This view is often presented under the motto that there is pragmatic encroachment on belief. The reasons that have been offered in favor of this view mirror the reasons that have been offered in favor of the idea that there is pragmatic encroachment on knowledge. I think those reasons aren't good for thinking there is such encroachment on knowledge (see e.g. Neta 2007; Reed 2010, 2012; Brown 2008, 2012; Lackey 2011; Roeber 2016). But I think they are even worse for thinking there is such encroachment on belief (see e.g. Cohen 2012; Ichikawa et al. 2012).

²² For example, Frankish (2004, 2009) and Ross and Schroeder (2014).

From these ideas, however, it follows—I think in a rather straightforward manner—that the requirements mentioned above do not impose rational strict liability and so are not genuine requirements of rationality.

The cases I go on to present are meant to illustrate this point. But the reason is rather simple and follows straightforwardly from what has been said so far: it is plausible to think that one dimension along which an agent's practical rationality can be assessed concerns the way in which her intentions relate to the space of possibilities that she regards as open for the purposes of deciding what to do, or, more generally stated, for the purposes of adopting, and maintaining, practical ends or goals ("for practical purposes," as I will put it from now on). I think this is an idea that is driving the coherentists, and I think it is a plausible idea. However, behind all the requirements stated above is the assumption that such a space is always carved up by the agent's beliefs, so that if an agent believes that p , then the possibility that *not* p is, for practical purposes, simply disregarded as an open possibility. But this is not true. For sure, beliefs oftentimes determine the space of possibilities that an agent regards as open, and so takes seriously, in practical deliberation. But sometimes, depending on what the odds are in a given deliberative context, they do not.²³ When they do not, it is implausible to think that they provide the framework in relation to which her practical rationality is to be assessed. The following cases try to illustrate this simple point.

3 Counter-examples to requirements of IB-coherence

3.1 Intention Detachment

According to Intention Detachment, rationality requires, roughly, that you do not intend to do something in case some condition obtains, believe that such condition obtains, believe that you will not do that thing unless you intend to do it, and not intend to do it.

Now consider the following case. I believe my dog has rabies. I am not yet sure. But I do believe he has it. I have already sent blood samples to the vet, and I am waiting to get official confirmation. Because of the risks of having a dog with rabies, and the suffering the dog himself would experience, I intend to kill him if he has rabies. Obviously, I love my dog, and I do not want to kill him. Given what is at stake, however, it seems plausible that I will not, and in fact should not, deliberate

²³ This discussion raises the broader question of when we should deliberate on the basis of beliefs and when on the basis of our credences. As I mentioned in footnote 15, given our cognitive limitations, we simply cannot do away with all-out beliefs altogether. Still, the question remains as to when we should deliberate on one rather than the other. I am afraid I do not have a theory to offer. My suspicion is that it ultimately depends on what maximizes expected utility. I would imagine that a *defeasible* disposition to deliberate on all-out beliefs is better in terms of expected utility than non-defeasible dispositions to deliberate only on subjective probabilities (given our cognitive limitations), or only on all-out beliefs (given the variability of stakes and odds). In any case, I hope I can defend my more limited thesis in this paper without committing to an overarching theory on this issue. I thank an anonymous referee for prompting me to address this point.

on the basis of my belief that he has rabies. I will, and should, rely on the probabilities I assign to his having rabies or not. Because I am relying not on my all-out belief but on my credences in my deliberation, and I still think there is a chance that my dog does not have rabies, I do not yet intend to kill him. But I do intend to kill him if he has rabies. Schematically, I intend that (*p if q*) and I believe that *q* (and also that, were I not to intend that *p*, then because of that *not p*; I will omit mentioning this last belief from now on). Yet, I don't intend that *p*. That I don't yet unconditionally intend to kill him makes a huge difference in my behaviour. I am not yet about to kill him, and I will not be until I am sure (or in any case sure enough) that he does have rabies.

So I fail to detach the consequent-intention from my conditional intention and my belief that the antecedent obtains. If the requirement of Intention Detachment imposes rational strict liability, I am guilty of practical irrationality. But, intuitively, I am guilty of no such thing. I am doing the most sensible thing I could do, given the situation. So the requirement of Intention Detachment does not impose rational strict liability and is not a genuine requirement of rationality.

One may want to object to this case by denying that the agent in question has the relevant belief or the relevant intention, so let me consider these objections in turn.

First, some people react to this example by claiming that it is implausible to think that, in a case like this one, I really believe that my dog has rabies. Instead, they claim that what I really believe is that he *most likely* has rabies (or something to that effect). Now, I do not want to reject that in such a scenario I may very well have this belief. But believing that my dog most likely has rabies doesn't prevent me from believing that he has rabies. We still need a reason to doubt that I have this belief. Some people say that the reason is that, in such a case, I believe the dog *might not* have rabies (or something to that effect), and that, if I believe this, then I don't really believe he has it. Again, though: the belief that the dog might not have rabies doesn't preclude the belief that he does. We still need a reason to think I cannot have this belief. Still others may want to claim that the fact that I am not, in that situation, willing to deliberate on the basis of the proposition that the dog has rabies is proof that I do not believe this. I have already explained why I reject this idea. In the end, I see no reason that would allow one to reject the case by simply denying that I could have this belief.

The second way in which one may want to reject this case is to say that it is implausible to think that, in such a scenario, I would really have the intention to kill the dog if he has rabies. Some people feel tempted to say that what I really intend is to kill the dog if (and perhaps only if) I am sure he has it, or maybe if (and only if) he has rabies and I am sure he does. Once again, I do not deny that it may be true of me that I have these intentions. Just as before, though, what I do want to deny is that having any of these intentions would in any way preclude my intending to kill the dog if he has rabies. It seems to me there are two states of affairs I am mainly concerned to avoid: the first one consists in my not killing the dog in the contingency that he has rabies, the other consists of my killing the dog in the contingency that he does not. My being determined to avoid these states of affairs is my intending, on the one hand, to kill the dog if he has rabies, and, on the other, to not kill him if it does not. It is because I have these intentions that I would form the

further intentions mentioned before. But those intentions are subsidiary: rather than preclude, they are formed precisely because I intend to kill the dog if he has rabies and not to kill him if he does not. So, again, I see no reason that would allow one to reject the case by simply denying that I could have the relevant intention.

For these reasons, I believe this case shows that the requirement of Intention Detachment does not impose rational strict liability and is not a genuine requirement of rationality.²⁴

3.2 Intention-Belief Consistency

Intention-Belief Consistency, roughly, requires that you do not intend to do what you believe you will not do. Some people think that this is in fact impossible.²⁵ I will not try to argue in favor of this possibility.²⁶ I assume it is possible. The question I want to ask is whether, assuming this is possible, it would necessarily be irrational. My suggestion is that it would not.

Consider: A fisherman, out at sea, hears by radio that a storm has formed and will soon hit him. He believes that the only way he has of making it out alive is reaching the shore before the storm reaches him. He is an experienced seaman, however, and, given the evidence, he correctly forms the belief that he will not be able to reach the shore in time. He is not completely certain of this. He knows there is a slight chance that he will make it back. But he is too experienced not to form the right conclusion in the face of the evidence. And the right conclusion is that the storm will hit him before he reaches the shore. Now, since he knows making it back is the only way he has of surviving, and since he thinks there is a slight chance of making it back, it is plausible to think that the fisherman would, and in fact should, form the intention to make it back.

The fisherman, then, would intend to do what he (justifiably) believes he will not do. If the requirement of Intention-Belief Consistency imposes rational strict liability, he is guilty of practical irrationality. But he is guilty of no such thing. He is doing the most sensible thing he could do, given the situation. So Intention-Belief

²⁴ As I mentioned in footnote 20, one could maintain that the odds of a given deliberative context can affect the amount of evidence needed for warranted belief. For example, one could claim that if it doesn't make rational sense in a given context to deliberate on the basis of a proposition, then it isn't rational in that context to believe it. I find this view implausible, but I cannot argue against it here. In any case, this view would imply that there *is* irrationality in the case I describe. The irrationality, however, would concern only the belief that my dog has rabies. It would not concern the way in which this belief relates to my conditional intention and the absence of an intention to kill the dog. In other words, my belief might be unwarranted, but there need not be any further, specifically structural irrationality in my (intending *p* if *q*, believing *q*, and not intending *p*). So, although this view does conflict with what I want to say about these cases (namely, that there is no irrationality involved), it does not conflict with my central claim (namely, that there is no structural irrationality associated with a violation of IB-coherence requirements). I thank an anonymous referee for prompting me to clarify this issue.

²⁵ Perhaps they believe this because they think an intention that *p* is, in fact, a belief that *p*. For a view like this see, for example, Setiya (2003, 2008, 2009), Velleman (1989, 2007) and Marusic and Schwenkler (2018).

²⁶ At least not directly, but what I go on to say will suggest why I think this view is false.

Consistency does not impose rational strict liability and is not a genuine requirement of rationality.

Some people will balk at this suggestion. They will insist that what makes sense for the fisherman to do in that situation is to *try* to make it back, and so perhaps, at most, to *intend to try* to make it back, but not to actually intend to make it back. This, they think, would be irrational. They will agree that it makes sense for the fisherman to *aim* to make it back, to adopt the *goal* or *end* to make it back, and to try as hard as he can to achieve it. But he better not actually intend it, because *that*—they think—would be irrational.

Just to be clear, then: nobody doubts that it would make perfect rational sense for the fisherman to make it his goal to reach the shore, and to try as hard as he can to do so. In fact, nobody doubts that it would make perfect rational sense for the fisherman to come to be disposed to think and act in the following manner:

- first, the goal of making it back is a settled, cross-temporally and counterfactually stable objective for him. (So, the agent represents this goal, monitors his progress towards it, and adjusts his behaviour so as to track it. Moreover, absent new reasons to reconsider, he will neither keep deliberating whether to pursue it, nor abandon it, he will stick to it as a settled object of pursuit);
- second, he is disposed to filter out as possible further objects of pursuit—as possible further *ends*—any state of affairs he believes to be inconsistent with his goal of making it back. (So, for example, if he believes that, were he to spend some time trying to catch more fish before heading back, then because of that he would not make it back, then he is disposed not to adopt the goal of spending any more time trying to catch more fish);
- third, he is disposed to figure out, and then to intend to take, any means he believes to be appropriate (implied, most desirable on balance, etc.) for his goal of making it back. (So, for example, if he believes that, were he not to start the motor right away, then because of that he would not make it back, he is disposed to intend to start the motor right away).

Notice, however, that these are, precisely, the dispositions that prominent philosophers of action tell us are characteristic of intention.²⁷ In other words, no one would deny that it makes perfect rational sense for the fisherman to come to be disposed towards the goal of making it back just as if he intended to make it back. But if this is so, then it would be obtuse to insist that it doesn't make rational sense for him to intend to make it back. If it makes rational sense for the fisherman to come to be disposed to think and act, and to actually think and act, exactly as if he

²⁷ See e.g. Bratman (1987), McCann (1986, 1989, 1991), Mele (1989, 1992), Adams (1986), Adams and Mele (1989) and Ferrero (2013). Following Bratman, we could add here that the agent is disposed to guide her own thought and action in accordance with her (perhaps implicit) acceptance of norms that demand corresponding norms of end stability, end consistency, and means-end coherence. Naturally, if one thinks that an intention that *p* metaphysically implies the belief that *p*, then one will think that an intention that *p* also involves all those dispositions characteristic of, or associated with, belief that *p*. Since I assume that intention does not imply belief, I assume that it is possible to intend that *p* when one is not disposed towards *p* in all the ways characteristic of, or associated with, belief that *p*.

intended to make it back, then it makes rational sense for him to intend to make it back.²⁸

Put differently: intentions, I assume, are realized by a certain cluster of dispositions. Such a cluster—action theorists tell us—involves dispositions to track the relevant end *e* and monitor progress towards it; not to deliberate further whether to pursue it, and to not drop it in the absence of reasons to reconsider. It includes a disposition not to adopt ends believed to be inconsistent with *e*.²⁹ And it involves a disposition to figure out and intend to take means believed to be appropriate to *e*. This dispositional profile realizes an intention that *e*.³⁰ This means that, if it isn't irrational for the fisherman to come to be so disposed towards the end of making it back, it isn't irrational for him to intend to make it back, since being disposed in that way realizes an intention to make it back.³¹ It isn't irrational for the fisherman to be

²⁸ I said the fisherman believes he will not make it back. I should add that he does not believe he will make it back. If this example is successful, then it also serves as a counterexample to a purported requirement that would demand that you do not (intend that *p* and not believe that *p*). Following Bratman (1987 p. 38), we could call such a requirement "Intention-Belief Completeness." I thank an anonymous referee for making me clarify this point.

²⁹ I will qualify this idea in footnote 30, to allow for the possibility of agents not deliberating, in specific occasions, on their all-out beliefs about relevant inconsistencies or appropriate means.

³⁰ Could intention not involve dispositions other than the ones I've identified? I suppose this is possible, but I do not know what these other dispositions could be. One natural suggestion is that an intention that *p* involves a disposition to make further planning on the assumption that *p* (e.g. Velleman (1989, 2007)). If intention did involve this disposition, then it *would* be irrational to come to be so disposed towards an end you believe will not obtain (or don't believe will obtain), since there does seem to be something rationally problematic about planning on the assumption that *p* when you actually believe that *not p* (or don't believe that *p*). Given what has been said already, though, it is easy to see why this disposition cannot be characteristic of intention in general. All we need to do is imagine a case where an agent intends *p* but where the risks of deliberating under the assumption that *p* outweigh the benefits of doing so. This may happen even in cases where the agent actually believes that *p*. Consider: a tightrope walker intends to cross the tightrope and believes (but is not certain) that he will do so. Conditional on his successfully crossing, not placing a safety net is the dominant strategy. Yet he does not deliberate whether to place the safety net or not on the assumption that he will be successful (after all, his life is at stake), but on the basis of the probabilities he assigns to his successfully crossing or not. Relative to such probabilities, the reasonable thing to do (we can assume) is to place the safety net. So he places the safety net. This shows that intending *p* does not necessarily involve a disposition to make further planning on the assumption that *p*. I do not know what else intending could involve. Take the fisherman: what else could his intending to get back, as opposed to his merely being disposed in the suggested manner, involve? I fail to see what this could be. In the absence of some definite proposal, and following what prominent theorists tell us about the dispositions that characterize intentions, I assume that the fisherman is disposed to get back in a way that realizes an intention to get back. I thank an anonymous referee for prompting me to address this issue.

³¹ Could one claim that what distinguishes an intention proper from the dispositional profile I've identified is *simply* that, when you believe that *not p* (or do not believe that *p*) it is irrational to intend *p*, but not to be disposed in the specified ways towards *p*? Well, this could not be the *only* distinguishing feature, since, I assume, rationality supervenes on the mind, so there cannot be a change in rational properties without a change in mental properties. Again, if intention involves dispositions different from the cluster I've identified, then there would be a difference in mental properties that could account for the difference in rational properties. I do not know what these different dispositions could be. As far as I can see, the mental state of the fisherman in my example is exactly like the mental state he would be in if he intended to make it back. So, if one is rational, the other one is too. I thank an anonymous referee for making me clarify this point.

so disposed. So Intention-Belief Consistency does not impose rational strict liability and is not a genuine requirement of rationality.

3.3 Intention Consistency

Intention Consistency requires, roughly, that you do not intend each of two ends you believe to be inconsistent.

Now consider: a doctor believes that substance X is the only way of treating life threatening condition C, but she believes that a dose of X of, or above, .5 ml would be lethal. A patient comes in with a critical case of C. The doctor believes that if the patient doesn't receive any treatment, she will die; that she has such an advanced case of C that no dose below .5 would be enough to save her; and that there is a very slight chance that a dose higher than .5 could save her. Supposing the doctor believes there are no further reasons (ethical, financial, legal, etc.) not to administer that dose, I think the doctor would, and should, intend to administer that dose. That is, given present odds, she would not, and in fact should not, deliberate on the basis of her belief that the dose is lethal. She would, and should, rely instead on the probabilities she assigns to it being lethal or not. And on the basis of the probabilities she assigns to the patient surviving conditional on her receiving that dose (extremely slight) versus the probabilities she assigns to the patient surviving conditional on her not receiving that dose (none), the reasonable course of action is to (intend to) administer that dose. But then it will be true of her that she intends to save the patient, all-out believes that giving her the dose will kill her, and yet intends to give her the dose.³² So she intends each of two ends she believes to be inconsistent. If the requirement of Intention Consistency imposes rational strict liability, then she is guilty of practical irrationality. But she is guilty of no such thing. She is doing the most sensible thing she could do in the situation. So the requirement of Intention Consistency does not impose rational strict liability and it is not a genuine requirement of rationality.

3.4 Means-End Coherence

Means-End Coherence requires, roughly, that you intend the means you believe to be implied by your ends.

Now consider: the zoo veterinarian intends that a newborn gorilla, Coco, survives. But Coco has a deficient organ. The vet thinks that the only thing that would save him is an organ transplant. But it so happens that the only possible donor is Coco's sibling, Mocha. The vet believes that taking the organ from Mocha would severely harm her. It wouldn't kill her, but it would certainly harm her. Since the

³² I said before that filtering out believed to be inconsistent states of affairs as possible goals is one of the characteristic features of intention. I am now suggesting that the doctor would not, and should not, filter out administering the dose, even though she believes this is inconsistent with saving the patient. Does this mean either that she would not, or should not, really intend to save the patient? I think it does not. What I think it shows is that we must re-describe the dispositions that characterize intentions to respect the fact that sometimes agents do not rely on (relevant) all-out beliefs in practical deliberation.

decision whether to make the transplant or not relies solely on the vet, she all-out believes that, were she not to intend that Coco receives the transplant from Mocha, then Coco would not survive. At the same time, she believes there is something else that could potentially save Coco: some kind of new, experimental treatment that promises to repair the relevant organ, but that hasn't been implemented with much success so far. Still, she thinks there is a very slight chance this could work. Suppose the vet has full authority to decide on the issue and put aside any possible ethical concerns. In this situation, I think she can intend that Coco survives and not intend that he receives the organ from Mocha. She might instead intend that he receives the experimental treatment. So this is a situation where she intends that Coco survives, all-out believes that, were she not to intend that he receives the transplant from Mocha, he would not survive, and yet does not intend that he receives that transplant. So she does not intend the means the intending of which she believes to be implied by her end. If the requirement of Means-End Coherence imposes rational strict liability, she is guilty of practical irrationality. But she is guilty of no such thing. She is doing the most sensible thing she could do, given the situation. So the requirement of Means-End Coherence does not impose rational strict liability and it is not a genuine requirement of rationality.

4 What conclusions should we draw?

I think coherentists should accept the previous cases. This does not mean that they should become myth theorists. It simply means that they need a more nuanced picture of when it is that all-out belief is relevant to determining an agent's practical rationality. What the previous cases show is the rather intuitive idea that the requirements of coherence that concern the presence or absence of intention in light of all-out belief are relevant to determining an agent's practical rationality only in cases where the agent regulates her intentions in light of the relevant all-out beliefs. They lose relevance when the agent regulates her intentions in light of the probabilities she assigns to the relevant contingencies, or in light of propositions she is, for whatever reason, accepting in a given context. This fact needs to be recognized by coherentists and reflected in their formulation of the relevant requirements.

Now, it may seem that coherentists should simply include in such formulations the condition that the agent deliberates in light of her relevant all-out beliefs. To take the example of Intention Consistency, it may seem that coherentists should simply say that rationality requires of an agent A that:

- *Intention Consistency Revised*: [if A believes, and deliberates under the assumption that, (p and q are inconsistent), then A does not (intend that p and intend that q)].

Once we formulate the requirement in this way, however, I think it becomes clear that including the condition that the agent believes the relevant proposition is misleading. After all, as I've argued, she could be rational while intending that p and

intending that q despite believing that (p and q are inconsistent) if this proposition did not, for whatever reason, play the role of carving up the space of possibilities that she regards as open for practical purposes. And, although I haven't argued for this idea here, it isn't implausible to think that she *would* be irrational (holding the intentions fixed) despite *not* believing this if that proposition did, nevertheless, and for whatever reason, play such a role. The crucial question, then, is whether that proposition plays the role of carving up the space of possibilities that she regards as open in her deliberation, not whether she believes it or not.

Perhaps, then, coherentists should simply say something roughly along the lines that rationality requires of an agent A that:

- *Intention Consistency Revised 2*: [if A deliberates under the assumption that (p and q are inconsistent), then A does not (intend that p and intend that q)].³³

Needless to say, this is still preliminary. Much more needs to be spelled out before we arrive at a satisfactory formulation of this requirement.³⁴ I will not attempt to do so here. In any case, I hope the suggestion is clear: coherentists should acknowledge the fact that beliefs do not always determine which possibilities we regard as open for practical purposes, and so do not always provide the framework in relation to which our practical rationality is to be assessed.

Naturally, coherentists who accept the cases I presented will still affirm, and myth theorists will still deny, that the appropriately re-formulated requirements are genuine.

Likewise, coherentists who accept these cases may still hold that requirements that are just like those of IB-coherence, except in that they are formulated in terms of certainty and not all-out belief, are genuine requirements of rationality. So, for example, coherentists who accept my case against Intention-Belief Consistency may still perfectly claim that rationality requires of an agent A that:

- *Intention-Certainty Consistency*: A does not [intend that p and be certain that *not* p].

Nothing I have said in this paper challenges this (or any similarly modified) requirement. Moreover, if certainty always carves up the space of possibilities that an agent regards as open for practical purposes—so that, necessarily, if an agent is

³³ Given the wide scope of the requirement, an agent could avoid irrationality by not intending that p , not intending that q , or by not deliberating under the assumption that (p and q are inconsistent). I thank an anonymous referee for suggesting that the condition that one deliberates under the relevant assumption should be within the scope of the requirement.

³⁴ For one thing, the requirement is now phrased partly in terms of the agent *doing* something (namely, deliberating under an assumption), rather than exclusively in terms of the agent having a certain set of attitudes. For another, it does not yet make explicit reference to a deliberative context. Following Bratman's (1992) terminology, perhaps we could say something roughly along the lines that an agent A "accepts" a proposition p in a deliberative context C, when p carves out the space of possibilities that A regards as open in C (where this involves, inter alia, a disposition to deliberate in C conditional on p). If so, then perhaps the requirement would say something along the lines that rationality requires of A that she does not [accept in C that (p and q are inconsistent), intend in C that p , and intend in C that q]. Again, much more needs to be spelled out. I will not try to do so here.

certain that p , then she deliberates under the assumption that p —then such requirements would be immune to the worries presented above.

Coherentists have also suggested that rationality requires of an agent A that:

- *Intention Non-Contradiction*: A does not [intend that p and intend that *not* p].

This is also a requirement of coherence, but it does not concern the presence or absence of intention in light of belief, so nothing I have said puts any pressure on it.

There is one more requirement I would like to consider, because, although it also concerns the presence of intention in light of belief, it appears to be, in an interesting sense, immune to the considerations presented so far. This is the requirement of Enkrasia. I turn to this issue now.

5 The case of Enkrasia

It is commonly thought that rationality requires of an agent A that:

- *Enkrasia*: A does not [believe (she ought that p), believe that (*not* p unless she intends that p), and not intend that p].

Enkrasia is a curious requirement in the landscape of rationality. To some people, it seems to enjoin something somehow different from simple coherence.³⁵ Be this as it may, if it is not a requirement of coherence proper, it is somewhere in the neighborhood of coherence, and it also concerns the presence or absence of intention in light of all-out belief. It would seem, then, that the kinds of considerations presented above should apply to it as well. That is, it would seem that, if an agent all-out believes, but is not certain, that she ought to do something (and holding her belief that she will not do it unless she intends to do it fixed, which from now on I will assume), then there could be cases where it would not be irrational for her not to intend to do it.

I myself had expected this, but I have come to see this as a mistake. To see why, let me consider a purported counterexample to Enkrasia, presented by Ralph Wedgwood (2013), that is supposed to exploit this very issue. It goes like this: you face two options, A and B, such that you must choose one and only one of them. You have a very high degree of confidence that you ought to do A, but you cannot rule out the possibility that you ought to do B instead. So you all-out believe, but are not sure, that you ought to do A. However, you are *certain* that, in the contingency that you are mistaken about what you ought to do, doing A would be catastrophic (it would involve, say, the destruction of the world), whereas in the contingency that you are right, doing B would only be slightly worse than doing A, and not really much of a problem. “In this case—Wedgwood says—it seems possible for you to be

³⁵ Andrew Reisner expresses concerns of this kind in his (2013). See also Broome (2013, pp 173–5).

rational, to have beliefs of this sort, and simultaneously to intend to do not A, but B instead.” (p. 491)

This, however, is not a counterexample to Enkrasia. This is because there are different senses of ‘ought,’ and the ought that figures in the cited belief is not the ought that figures in the requirement. The ought of Enkrasia is not the sometimes called ‘objective’ ought of full information. It is the so called ‘subjective’, ‘practical’ or ‘deliberative’ ought. Since we oftentimes deliberate without full information, these oughts can, and oftentimes do, diverge. As different theorists have argued, however, it isn’t necessarily irrational not to intend to do what you believe you ought objectively to do.³⁶ This is a point that defenders of the requirement of Enkrasia recognize.³⁷ So this cannot be the ought they have in mind when they defend Enkrasia.

As Broome (2013, pp 24–25) emphasizes, the ought of Enkrasia is identified, *precisely*, as the ought of which it is true that, if you believe you ought, in that sense, to ϕ , then you are irrational if you do not intend to ϕ . This is not the ought that figures in the belief that you ought to A in Wedgwood’s example.

Wedgwood tells us you believe you ought to A. Well, perhaps you believe that you ought objectively to A, but given that you believe there are huge risks associated with A-ing in case you are mistaken, and only negligible costs associated with B-ing in case you are not, it is not the ought that would guide your action if you were rational. Since it is not the action-guiding ought of deliberation, you don’t believe that you ought to A in the sense that figures in Enkrasia, and this is not a counterexample to it. In fact, if you are rational, besides believing that you ought, objectively, to A, you will believe that you ought, in light of the information you *do* have, to B (after all, if you A you risk destroying the whole world!). This, then, is the ought that, if you are rational, guides your action. It is the practical or deliberative ought. This is the ought of Enkrasia.

Now, there is an interesting question of what exactly accounts for this peculiarity of Enkrasia. I suspect the reason is that—at least as far as I can see—every reason that would count against doing what would be best conditional on the truth of the belief that you ought—in the relevant sense—to do something, works by way of counting as a reason to doubt its truth, and so against holding such a belief in the first place. This is not the case with all the previous requirements. With respect to them, there may be plenty of reasons not to do what would be best conditional on the truth of the relevant belief that do not, in any way, constitute evidence against its truth. Such reasons are normally considerations that point to the risks of ignoring the possibility that the belief might be false, without putting into question its truth. For example, that the fisherman’s life is at stake counts against his ignoring the

³⁶ Curiously, a case that is commonly used to illustrate this point is structurally identical to a second purported counterexample to Enkrasia that Wedgwood presents. The case is considered, among others, by Regan (1980), Jackson (1991), Ross (2006), Kiesewetter (2011) and Finlay (2014). None of them take it to be a counterexample to Enkrasia. In fact, they all take it to show, precisely, that there are cases where you might believe (and even be certain) that you ought, objectively, to do something you ought not, subjectively, to do.

³⁷ e.g., Broome (2013, ch 3).

possibility of making it back, without constituting evidence that his belief that he will not make it back is false. The same is true of all the other requirements.

In the case of Enkrasia, though, considerations that point to the risks of doing what would be best conditional on the truth of the belief that you ought—in the relevant sense—to do something, work—as far as I can see—by way of pointing to the possible bad features or consequences of doing that thing. But such considerations, in turn, constitute evidence against the truth of the belief, because what you ought to do is in part determined by the risks and possible bad features or consequences of doing it. To return to Wedgwood's example: the consideration that, if you A in the contingency of being mistaken about what you ought to do, you will destroy the whole world, counts as a reason to doubt that you really ought—in the relevant sense—to A. It counts as evidence that you ought, instead, to take the much less risky option of doing B. This feature is unique to Enkrasia, and—as far as I can see—it is why Enkrasia isn't affected by the considerations presented above.³⁸

6 Conclusion

It is natural to think that rationality requires that we somehow regulate our intentions in relation to the space of possibilities that we regard as open for practical purposes. Coherentists are plausibly right to think this. But it is a mistake to assume that such a space is always carved up by the agent's beliefs. Coherentists have been wrong to assume this. Because of this, (most of) the requirements of intention in light of belief that they have tried to elucidate and justify are not genuine requirements of rationality. Still, it might perfectly well be the case that, for each of those purported requirements, there is a corresponding genuine requirement that is formulated not in terms of what the agent believes, but in terms of the propositions that carve up the space of possibilities that she regards as open for practical purposes. Such requirements would still count as requirements of coherence.

The cases I presented depend on a view about the nature of all-out belief that is supported by what I take to be two extremely plausible theses: The first is that one shouldn't be willing to bet at any odds on the truth of p if one is not certain that p ; the second is that belief does not imply certainty. The view they support is that one shouldn't be willing to deliberate, in every context and regardless of the odds, on the basis of one's (relevant) beliefs.

All the cases I presented depend on this view about all-out belief. Now, although I offered some of the reasons why I think this view is extremely plausible, I did not exactly try to give a proper defense of it. So if one is thoroughly convinced that the

³⁸ In other words: suppose you believe, but are uncertain, that you ought—in the relevant sense—to A. What could make your refraining from intending to A, given your belief, not irrational? As far as I can see, only your believing that it would be better—in the relevant sense—to refrain from A-ing (because of the risks associated with A-ing). But then you do not believe that you ought, in the relevant sense, to A. This is not to say that there might not be *different* reasons (not having to do with uncertainty and risk) to doubt that Enkrasia is a genuine requirement of rationality, as several philosophers have argued (e.g. Audi 1990; Arpaly 2003; Brunero 2013).

requirements of intention in light of belief defended by IB-Coherentists are genuine, one might perfectly well use my own cases to reject the view of all-out belief on which they depend. As they say, one philosopher's modus ponens is another's philosopher's modus tollens. So holding these requirements as fixed theoretical points would provide a reason for a radically revisionist view of the nature of belief. I imagine some people will welcome this result. I myself think we should stick to the traditional view of belief, and instead revise our understanding of exactly how, and when, rationality requires that we regulate our intentions in light of our beliefs.

Acknowledgements Open access funding provided by University of Vienna. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No 740922). Thanks to Facundo Alonso, Sam Asarnow, Olle Blomberg, Michael Bratman, John Broome, Rahul Chaudhri, David Hills, Nick Hoff, Javier Peñalosa, David Taylor, a very helpful anonymous referee for *Philosophical Studies*, participants in the *Varieties of Agency* workshop at Stanford, members of the ERC Project "The Normative and Moral Foundations of Group Agency," and audiences at the University of Minnesota and the 10th congress of the German Society for Analytic Philosophy, for their valuable feedback.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Adams, F. (1986). Intention and intentional action: The simple view. *Mind & Language*, 1(4), 281–301.
- Adams, F. & Mele, A. (1989). The role of intention in intentional action. *Canadian Journal of Philosophy*, 19(4), 511–531.
- Alonso, F. M. (2014). What is reliance? *Canadian Journal of Philosophy*, 44(2), 163–183.
- Alonso, F. M. (2016). Reasons for Reliance. *Ethics*, 126(2), 311–338.
- Arpaly, N. (2003). *Unprincipled virtue*. Oxford: Oxford University Press.
- Audi, R. (1990). Weakness of will and rational action. *Australasian Journal of Philosophy*, 68(3), 270–281.
- Bratman, M. E. (1987). *Intention, plans, and practical reason*. Cambridge: Harvard University Press.
- Bratman, M. E. (1992). Practical reasoning and acceptance in a context. *Mind*, 101(401), 1–16.
- Broome, J. (2013). *Rationality through reasoning*. Hoboken: Wiley-Blackwell.
- Brown, J. (2008). Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2), 167–189.
- Brown, J. (2012). Practical reasoning, decision theory and anti-intellectualism. *Episteme*, 9(1), 43–62.
- Brunero, J. (2013). Rational Akrasia. *Organon F*, 20(4), 546–566.
- Buchak, L. (2013). Belief, credence, and norms. *Philosophical Studies*, 169(2), 285–311.
- Christensen, D. (2007). *Putting logic in its place*. Oxford: Oxford University Press.
- Clarke, R. (2013). Belief Is Credence One (in Context). *Philosophers' Imprint* 13:1–18.
- Cohen, S. (2012). Does practical rationality constrain epistemic rationality? *Philosophy and Phenomenological Research*, 85(2), 447–455.
- de Finetti, B. (1970). *Theory of probability*. Hoboken: Wiley.
- Dodd, D. (2017). Belief and certainty. *Synthese*, 194(11), 4597–4621.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.

- Ferrero, L. (2013). Can I only intend my own actions? In D. Shoemaker (Ed.), *Oxford studies in agency and responsibility* (Vol. 1, pp. 70–94). Oxford: Oxford University Press.
- Finlay, S. (2014). *Confusion of tongues: A theory of normative language*. New York: Oxford University Press.
- Frankish, K. (2004). *Mind and supermind*. Cambridge: Cambridge University Press.
- Frankish, K. (2009). Partial belief and flat-out belief. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of belief* (Vol. 342). Berlin: Springer.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441–458.
- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is weak. *Philosophical Studies*, 173(5), 1393–1404.
- Ichikawa, J. J., Jarvis, B., & Rubin, K. (2012). Pragmatic encroachment and belief-desire psychology. *Analytic Philosophy*, 53(4), 327–343.
- Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461–482.
- Joyce, J. M. (1999). *The foundations of causal decision theory* (Vol. 112). Cambridge: Cambridge University Press.
- Kaplan, M. (1998). *Decision theory as philosophy*. Cambridge: Cambridge University Press.
- Kiesewetter, B. (2011). “Ought” and the perspective of the agent. *Journal of Ethics & Social Philosophy*, 5(3), 1–24.
- Kiesewetter, B. (2017). *The normativity of rationality*. Oxford: Oxford University Press.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(455), 509–563.
- Kolodny, N. (2007). How does coherence matter? *Proceedings of the Aristotelian Society*, 107(pt 3), 229–263.
- Kolodny, N. (2008a). The myth of practical consistency. *European Journal of Philosophy*, 16(3), 366–402.
- Kolodny, N. (2008b). Why be disposed to be coherent? *Ethics*, 118(3), 437–463.
- Lackey, J. (2011). Acting on knowledge. *Philosophical Perspectives*, 24(1), 361–382.
- Levi, I. (1991). *The fixation of belief and its undoing: Changing beliefs through inquiry*. Cambridge: Cambridge University Press.
- Maher, P. (1986). The irrelevance of belief to rational action. *Erkenntnis*, 24(3), 363–384.
- Maher, P. (1993). *Betting on theories*. Cambridge: Cambridge University Press.
- Marusic, B., & Schwenkler, J. (2018). Intending is believing: A defense of strong cognitivism. *Analytic Philosophy*, 59(3), 309–340.
- McCann, H. J. (1986). Rationality and the range of intention. *Midwest Studies in Philosophy*, 10(1), 191–211.
- McCann, H. J. (1989). Intending and planning: A reply to Mele. *Philosophical Studies*, 55, 107–110.
- McCann, H. J. (1991). Settled objectives and rational constraints. *American Philosophical Quarterly*, 28, 25–36.
- Mele, A. R. (1989). Intention, belief, and intentional action. *American Philosophical Quarterly*, 26(1), 19–30.
- Mele, A. R. (1992). *Springs of action: Understanding intentional behavior* (Vol. 103). Oxford: Oxford University Press.
- Neta, R. (2007). Anti-intellectualism and the knowledge-action principle. *Philosophy and Phenomenological Research*, 75(1), 180–187.
- Raz, J. (2005). The myth of instrumental rationality. *Journal of Ethics & Social Philosophy*, 1(1), 1–28.
- Reed, B. (2010). A defense of stable invariantism. *Noûs*, 44(2), 224–244.
- Reed, B. (2012). Resisting encroachment. *Philosophy and Phenomenological Research*, 85(2), 465–472.
- Regan, D. H. (1980). *Utilitarianism and co-operation* (Vol. 33). Oxford: Oxford University Press.
- Reisner, A. (2013). Is the Enkratic principle a requirement of rationality? *Organon F*, 20(4), 437–463.
- Roeber, B. (2016). The pragmatic encroachment debate. *Noûs*, 52(1), 171–195.
- Ross, J. (2006). Acceptance and practical reason. *PhD Dissertation*. Rutgers University.
- Ross, J., & Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*, 88(2), 259–288.
- Savage, L. (1972). *The foundations of statistics* (2nd ed.). New York: Dover Publishing.
- Setiya, K. (2003). Explaining action. *Philosophical Review*, 112(3), 339–393.
- Setiya, K. (2008). Practical knowledge. *Ethics*, 118(3), 388–409.
- Setiya, K. (2009). Practical knowledge revisited. *Ethics*, 120(1), 128–137.

- Shah, N. (2003). How truth governs belief. *The Philosophical Review*, 112(4), 447–482.
- Shah, N., & Velleman, J. D. (2005). Doxastic deliberation. *The Philosophical Review*, 114(4), 497–534.
- Velleman, J. D. (1989). *Practical reflection* (Vol. 94). Princeton: Princeton University Press.
- Velleman, J. D. (2000). On the aim of belief. In *The possibility of practical reason* (pp. 244–281). Oxford: Oxford University Press.
- Velleman, J. D. (2007). What good is a will? In A. Leist & H. Baumann (Eds.), *Action in context*. Berlin: de Gruyter/Mouton.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443.
- Wedgwood, R. (2002). The aim of belief. *Noûs*, 36, 267–297.
- Wedgwood, R. (2013). Akrasia and uncertainty. *Organon F*, 20(4), 484–506.
- Weisberg, J. (2009). Commutativity or holism? A dilemma for conditionalizers. *The British Journal for the Philosophy of Science*, 60(4), 793–812.
- Worsnip, A. (2016). Belief, credence, and the preface paradox. *Australasian Journal of Philosophy*, 94(3), 549–562.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.