

Faith in Kant

Guy Longworth¹
University of Warwick
2.9.15

1. Introduction: faith and trust. Can it ever be reasonable to trust other people, or to rely upon them, without evidence of their trustworthiness or reliability? I shall be exploring an affirmative answer to that question drawn from the work of Immanuel Kant. According to Kant, there can be reasons for holding things true which are *practical* rather than theoretical or evidential. There can therefore be practical reasons for holding true that other people are trustworthy or reliable. I shall explain Kant's view about holding things true for practical purposes and make a start at assessing it.

As limited beings, we often have cause to rely upon the acts and attitudes of other limited beings. In order for it to be reasonable for us to rely upon others, it seems that we must have reasons to hold that those we rely upon will be appropriately reliable. And yet it seems that many of the circumstances in which we are required to rely upon others are precisely circumstances in which we lack sufficient evidence to hold that they will be reliable. Thus, our practical and epistemic limitations conspire to produce a puzzle: how, if at all, can it be reasonable for us to do what it seems that we must: rely upon, or trust, other limited beings?

Different cases might well call for different responses. One response to a case of the puzzle would be to withhold trust: we might sacrifice the benefits to be secured by relying on others in order to preserve our claim to reasonableness. A second response would be to rely upon others whilst acknowledging that our doing so makes us unreasonable. A third response would be to attempt to argue that our epistemic limitations are less stark than the puzzle makes them out to be—that, in a particular case, we have sufficient evidence to hold reasonably that those that we need to rely upon will be reliable. A fourth response would be to attempt to show that our reliance on others can be reasonable despite our lacking evidence, because we can possess non-evidential reasons that are sufficient for holding reasonably that others will be reliable. (Limiting cases of the third and fourth responses might be grounded on the view that our practice of relying on others can be reasonable in the absence of reasons.)

¹ I'm grateful for discussion and comments to Tom Crowther, Naomi Eilan, Paul Faulkner, Giulia Felappi, Keith Hossack, Nils Kürbis, Hemdat Lerman, Daniel Morgan, Johannes Roesler, Barry C. Smith, Matthew Soteriou, and Mark Textor. I'm also grateful to audiences at King's College London, the London Institute of Philosophy, and Liverpool University for discussion.

I shall be exploring the prospects of a particular version of the fourth response to the puzzle. Kant's aim, in the discussions on which I'll focus, is to secure the reasonableness of what he calls *moral faith*. Moral faith, for Kant, is a form of holding something true that is grounded in practical reason. In particular, it is grounded in the moral demands to which, according to Kant, practical reason is subject. Thus, moral faith is grounded in practical rather than theoretical reason. Specifically, moral faith is a matter of holding, on moral grounds, that we are free, that we are immortal, and that God exists.

Kant himself understands faith in God to involve a form of trust:

Faith (simply so called) is trust in the attainment of an aim the promotion of which is a duty but the possibility of the realization of which it is not possible for us *to have insight into...* (Kant 1793: 5: 472)

Faith means the confidence that, so long as we have done everything possible to us, God will supply what does not lie in our power... The sole object of spiritual trust is in the pure morality, the holiness of man, and then his eternal blessedness under the conditions of morality... Hence, in order that our trust may coincide with the plan of wisdom, it must be a wise trust, and unconditional, so that we believe in general that God, in His goodness and holiness, will both lend us His aid in regard to acting morally, and also allow us to participate in blessedness. (Kant 1784–5: 27: 321–2)

However, my central interest in Kant's discussion is not in his entitlement to the specific objects of faith that he emphasizes. My central interest is rather in the principles that Kant employs in underwriting faith in those objects and, then, in the possibility of using those principles, or principles like them, in underwriting trust in other people.

I shall proceed as follows. In section 2 (*Kant on moral faith*), I'll expound Kant's case for moral faith, specifically in God, and I'll highlight two central principles that figure in that case. I'll pursue the standing of those principles in sections 3 to 5. Section 3 (*Problematizing faith*) raises the central difficulty to be pursued, the need to divorce the attitude involved in faith from evidentially grounded belief. Sections 4 (*Defending Kant's principles*) and 5 (*Faith and evidence*) develop a case for the required divorce. In the concluding section 6, I'll briefly consider the prospects of extending Kant's case for moral faith in order to underwrite a form of trust, or faith, in other people, and raise two questions about that extension.

2. *Kant on moral faith*. Let's begin with a sketch of Kant's case for moral faith in God. This will enable us to see the main moving parts of Kant's case, and thus to isolate which, amongst those moving parts, are of greatest relevance for our more general interest in the reasonableness of extra-evidential trust. (Kant presents versions of his basic case for moral faith in a large number of places, in particular: 1781/1787: A795/B823–A831/B859; 1786; 1788: 5: 107–148; 1790: 5: 397–415, 5:434–484; 1793: 6: 3–11. The level of detail at which I'll be working excuses me

from detailed engagement with these texts and from taking into account the fine differences between them.)

Kant's aim is to show that it is reasonable to hold true, with certainty, that there exists a being that is both willing and able to organize the world in accord with a moral design. Kant takes himself to have demonstrated, in the course of the first critique, that it is impossible for us to *know* that there is (or that there isn't) such a being. It is impossible for us to have purely theoretical, or evidential, grounds for holding with certainty that such a being exists. For evidence that is accessible to us must be grounded in our intuition, and our intuition is restricted to appearances; but if God exists, then He is super-sensible, so evidence for or against His existence is not to be found amongst the appearances. So, any reasons for holding with certainty that God exists must be non-evidential.

Such non-evidential reasons are provided, according to Kant, by practical reason. For practical reason gives rise to the moral law. Practical reason demands, moreover, that we act from respect for the moral law. And Kant argues that we can act rationally from respect for the moral law only on condition that we set what he calls *the highest good* as our ultimate end. The highest good unifies all more specific moral goods. It would be a state of the world in which all persons acted only from respect for the moral law and, moreover, one in which their doing so gives rise to their also being proportionately (so maximally) happy. But, according to Kant, the rationality of our setting the highest good as our ultimate end depends upon its being reasonable for us to hold that the highest good is in principle attainable. And its being reasonable for us to hold that the highest good is in principle attainable requires that it is reasonable for us to hold true anything on which its in principle attainability depends. In particular, it requires that it is reasonable for us to hold two things. First, it requires that it is reasonable for each of us to hold not only that we will act only from respect for the moral law, but in addition that everyone else will do so too. And second, it requires that it is reasonable for us to hold that if everyone acts only from respect for the moral law, then their doing so will be efficacious in determining that everyone is made happy in proportion to their moral worthiness. The former requirement seems to require reasonable faith or trust in one's own potential morality as well as in that of all other people. (Kant 1781/1787: A810/B838, 1785: 106.) With respect to the first requirement, there appear to Kant to be only two options: either it must be reasonable to hold that nature is organized so that morality leads to proportionate happiness, or it must be reasonable to hold that the hedonic efficacy of morality is enforced supernaturally, by God. But it is not reasonable—we lack evidence for holding—that nature is so organized. Hence, the reasonableness of our setting the highest good as our end depends upon the reasonableness of our holding that God will enforce the efficacy of morality in determining happiness. So, since it is reasonable for us to act from respect for the moral law, it is reasonable for us to set the highest good as our end; and since it is reasonable for us to set the highest good as our end, it is reasonable for us to hold that God exists and will ensure that everyone's doing their part—everyone's acting from respect for the moral law—will lead to their proportionate happiness.

The sketched argument requires development at a number of points.

First, it needs to be explained why we should hold that it is reasonable to act from respect for the moral law (at least, *pro tanto*, on the assumption that the demands imposed by reason are consistent).

Second, it needs to be explained why acting from respect for the moral law demands that we set the highest good, as conceived by Kant, as our end. Kant himself vacillates over the precise relationship between the moral law and the highest good and, indeed, about the precise nature of the highest good. And his most detailed accounts of the relationship fail to make transparent how precisely the rational demand that we act from respect for the moral law transmits to a demand that we set the highest good as our ultimate end.

Third, it needs to be explained why it can be reasonable to set an end only on condition that it is also reasonable to hold that the end is in principle attainable. For the argument would fail, not only if it was reasonable to set an end that one held to be unattainable, but also if it were reasonable to set an end that one reasonably didn't hold to be unattainable. For, as noted, the argument of the first critique precludes not only our acquiring evidence for the existence of God, but also our acquiring evidence against His existence. So, we can have no evidential grounds for holding that God doesn't exist, and so no evidential grounds for holding that the highest good is unattainable. Thus, if the only requirement on rationally setting an end were that we lack reason to hold that it is unattainable, we might be in that position independently of our possessing positive reasons to believe that God exists.

Fourth, it needs to be explained what precludes our holding that nature itself is organized in accord with the attainability of the highest good. For even if we agree with Kant that we currently lack evidence that it is so organized, the absence of such evidence need not be due to a failure of correlation between moral worthiness and happiness. It might be due, instead, to the current paucity of morally worthy people. Of course, if we were required to hold positively that the highest good is attainable in the absence of such evidence, our doing so would require a form of moral faith. But it would require faith concerning the organization of nature, rather than the offices of a supernatural organizer.

The third and fourth issues are connected. For suppose that rationally setting an end required only that one lacks reasons for holding that the end is unattainable and not, in addition, that one possesses reasons for holding that it is attainable. In that case, as noted above, the absence of evidence that nature is organized so as to sustain the attainment of the highest good would leave open that we may nonetheless set the highest good as our end. And so even if practical reason required us to set the highest good as our end, it wouldn't, in addition, require that we held that nature was organized—intrinsically, or by God—so as to ensure the sustainability of that end. However, an attempt might be made to develop Kant's argument by appeal, not to the absence of evidence that nature is organized so as to sustain the attainment of the highest good, but rather to the presence of evidence that nature is organized so as to preclude its attainment. In order to make that case, more would be required than merely to point to the ubiquity of unhappiness, even amongst the most morally worthy. For, first, one would need to show that they were less happy than they should be, given their moral worth. And, second, one would need to provide evidence that the failure of

appropriate correlation between worthiness and happiness is of a type liable to be preserved into a circumstance in which all persons are morally worthy.

Fifth, it needs to be explained why the only alternative to the attainability of the highest good being secured by the intrinsic organization of nature is that it is secured by the organizational work of God.

Finally, sixth, it needs to be explained how the *pro tanto* practical reasonableness of our acting from respect for the moral law can transmit so as to make reasonable our holding true that God exists, rather than imposing on us a requirement to obtain independent theoretical grounds for holding that God exists. Now the requirement that we find independent grounds might be one that we cannot in principle satisfy. In that case, and assuming that the remainder of the argument were cogent, we would be presented with a conflict between the demands set by practical reason and demands set by theoretical reason. That might be, to say the least, unfortunate. But it's not clear why the hope of avoiding such an outcome should be taken to provide reasons for treating as spurious the apparent demands either of practical or of theoretical reason. And even if it did provide such reasons, we would then have to confront the further question, why should it be the apparent demands of theoretical reason that are to be surrendered, rather than those of practical reason, and with them, the demand that we act from respect for the moral law?

Clearly, the proper assessment of Kant's argument would require extended labour, and would need to go to the heart of his moral theory. (An excellent attempt is made in Wood 1970. See also Gardner 2006; Kleingeld 1998; Neiman 1994; Timmermann 2009; Watkins 2010; Williaschek 2010.) However, for present purposes, I'll focus on the need to develop the sketch at two main points, the third and sixth of the points just listed: the purported demand on the rational setting of ends, according to which it requires reason to hold that those ends are in principle attainable; and the claim that that demand can serve to furnish extra-evidential reasons for holding things true. In doing so, I'll be pursuing three questions that can only be addressed separately with some artificiality:

Q1. What cognitive requirements must one meet if one is rationally to set an end? Must one, for example, hold true that the end is in principle attainable?

Q2. What is the nature of the attitude of holding true that figures in meeting those cognitive requirements? Is it a paradigmatic form of belief, or does it take some other form?

Q3. To what extent can the required form of holding true be sustained rationally by the needs of practical reason? Does the answer to that question depend on whether the operative needs arise from an absolute demand on practical reason, that we follow the moral law? And does it depend on whether the required form of holding true concerns only things about which we cannot know?

3. *Problematizing faith.* In order to fix ideas, it will be useful to sketch an argument that begins from Kant's principle governing the rational setting of ends and terminates in the conclusion that there is an evidential requirement on setting ends that precludes the possibility of moral faith.

The argument begins from a principle that Kant accepts:

(P1) One can rationally set an end, *E*, only if one rationally holds true that *E* is in principle attainable.

The attainability in principle of *E* typically will depend, and typically will be known to depend, upon two sorts of factors. First, it will depend upon what agents can do towards bringing the end about, conditional on their contingent circumstances. That is, it will depend upon the range of things that are within an agent's power to bring about, given specific ways things outside their power are. Second, it will depend upon those aspects of the agent's contingent circumstances that are outside their power to affect. Since the successful attainment of one's ends is often dependent upon the activities of other agents, the latter conditions will often include those activities, and the willingness on which they depend. Let's label the former type of necessary condition on the attainability of *E* *agential conditions*, and the later type of necessary conditions *environmental conditions*. A consequence of (P1), then, is (P2):

(P2) One can rationally set an end, *E*, only if, with respect to all conditions that one holds to be either agential conditions or environmental conditions, one rationally holds true that those conditions obtain or will obtain.

That principle now interacts with principles purportedly governing what one may rationally hold true so as to generate the putative difficulty for Kant's account. So, suppose that, in addition to (P2), one held (P3) and (P4):

(P3) One can rationally hold true that a condition will obtain only if one can rationally believe that the condition will obtain.

(P4) One can rationally believe that a condition will obtain only if one possesses evidentiary support for the obtaining of the condition.

An attempt to motivate principles (P3) and (P4) might make appeal to a conception of the constitutive aims of holding true on which the constitutive aim of holding something true is truth or knowledge. Plausibly, such a conception of holding true would align the constitutive aim of holding true with the constitutive aim of belief, so sustaining (P3). And it's plausible that it would, thus, sustain (P4) by sustaining the evidentiary requirement in (P4). On the basis of those principles, we could derive an evidentiary requirement on the setting of ends to the effect that one can rationally set an end only if one possesses evidentiary support for the obtaining of agential and environmental conditions on the attainability of the end.

Such a requirement on the rational setting of ends might appear too strong. For, at least from an agent's own perspective, answering the question whether

agential conditions will obtain seems to depend, not on the acquisition of evidence, but rather on making up one's mind on the ends that they will pursue. However, we can finesse that concern, for present purposes, by stipulating that, with respect to agential conditions, the evidentiary requirement can be met on the basis of so setting an end. That leaves the environmental conditions. It would be natural to hold that evidence for the obtaining of environmental conditions would have to be decision-independent evidence.

On that basis, our principles dictate that, with respect to ends for which there are non-agential environmental conditions, one can rationally set an end, *E*, only if one possesses evidence that *E*'s environmental conditions obtain. Thus, in cases in which the environmental conditions include the activities of other agents, one must possess evidence that the agents will be willing and competent to undertake those activities. The conclusion of Kant's argument for moral faith was that one can rationally set ends with environmental conditions for the obtaining of which one doesn't possess—indeed, for which one cannot possess—any evidence at all. So, Kant's claims for moral faith conflict with the argument we have just developed, and it is incumbent on the defender of Kant to say something in response. Minimally, it seems, they must provide grounds for rejecting either (P3) or (P4). In doing so, they should say something about the constitutive aims of holding true for practical purposes that supports the rejection of those principles. That is a large project. I shall make a start on pursuing it in the following two sections, by developing an alternative conception of holding true for practical purposes. According to the alternative conception, the operative form of holding true is not a paradigmatic form of belief (so that (P3) is unsupported) and, so, is not subject to the same evidentiary demands as are paradigmatic forms of belief (so that (P4) is unsupported).

4. *Defending Kant's principles.* Kant's argument depends upon the following principles, purportedly governing the connection between one's rationally setting an end and one's holding true that the end is in principle attainable:

(P1) One can rationally set an end, *E*, only if one rationally holds true that *E* is in principle attainable.

(P2) One can rationally set an end, *E*, only if, with respect to all conditions that one holds to be either agential conditions or environmental conditions, one rationally holds true that those conditions obtain or will obtain.

Setting an end for oneself is freely adopting the aim of doing what one can in order to bring about, or preserve, the end. It has as upshot its passive counterpart, having an end, itself a mode of intention. The question whether (P1) and (P2) are correct is, therefore, a near relative of the question whether (rationally) intending to ϕ entails (rationally) believing that one will ϕ . I'll approach the former question indirectly, via its connections with the latter.

In attempting to address the question whether intending to ϕ entails believing that one will ϕ , it is important to avoid conflating distinct issues. It is especially important to avoid two tempting conflations.

First, it is important to avoid conflating the question whether a rational intention to ϕ entails holding true that one will ϕ with the distinct question whether a rational intention to ϕ entails believing that one will ϕ . Although it may be that all cases of holding true are cases of believing, it shouldn't be assumed from the outset that they are. Reasons to think that intending doesn't entail believing may be precisely reasons to think that holding true doesn't entail believing.

Second, it is important to hold distinct two forms that may be taken by intention: on the one hand, having (or being) decided to ϕ and, on the other hand, merely having (set) the end of ϕ -ing. If I have (or am) decided to ϕ , then I have the end of ϕ -ing. However, the converse entailment fails. For I can have, and act upon, ends that I can't bring about merely by so acting—for example, I can have the end of winning a fair lottery. By contrast, I cannot decide to win a fair lottery; at best, I can decide to try to win it. It would be natural, therefore, to expect the cognitive requirements on deciding to ϕ to differ from those on deciding on the end of ϕ -ing. Suppose, for example, that deciding to ϕ entails believing that one will ϕ . And suppose, furthermore, that deciding on the end of ϕ -ing entails deciding to try to ϕ . It would follow that deciding on the end of ϕ -ing entails believing that one will try to ϕ . But it would be implausible to hold, on those grounds, that believing that one will try to ϕ entails believing that one will ϕ . Moreover, given that thought and talk about what someone intends can target either what they have decided to do or their wider ends—e.g., what they have decided to try to do—we would naturally expect our judgments about the cognitive requirements on intending to be pulled in opposing directions. On the one hand, insofar as our attention is focused on the requirements on deciding to ϕ , we will find it more intuitive that they include holding that one will ϕ . On the other hand, insofar as our attention is directed onto the requirements on merely deciding on the end of ϕ -ing, without deciding to ϕ , we will find it less intuitive that they include holding that one will ϕ , and will expect to discern, instead, somewhat weaker requirements on the rational setting of ends.

Although it is possible to decide on an end that depends upon one's ϕ -ing without thereby deciding to ϕ , we can nonetheless achieve insight into deciding on an end via reflection on deciding to ϕ . Let's begin, then, by considering principle (P5), governing rationally deciding to ϕ :

(P5) One can rationally decide to ϕ only if one rationally holds true that one will ϕ .

Two central questions that arise with respect to (P5) are the following. First, what reason is there to accept (P5) as a requirement on rational decision, in favour of weaker principles? Second, how should we understand the nature of the attitude of holding true that is embedded in (P5)? How, in particular, does that attitude relate to evidence that one will ϕ , and in particular evidence that obtains

independently of one's decision to ϕ ? And how, if at all, does one's holding, in that way, that one will ϕ figure in bringing it about that one will?

One reason for endorsing (P5) is that it seems obviously correct. Going beyond that, it can figure in explaining some of the other rational norms to which deciding to ϕ is subject. For example, it seems plausible to accept that any reasonable account of holding true will be subject to a requirement like the following:

(P6) One can rationally hold that p , and hold that its being true that p entails its not being true that q , only if one doesn't hold that q .

On the basis of (P5) and (P6), we can explain why one can't rationally both decide to ϕ and decide to ψ while one holds that its being true that one will ϕ entails its not being true that one will ψ . For according to (P5), one's rationally deciding to ϕ entails that one rationally holds that one will ϕ , and one's rationally deciding to ψ entails that one rationally holds that one will ψ . But according to (P6), one can't rationally hold that its being true that one will ϕ entails that one will not ψ and at the same time hold that one will ϕ and that one will ψ . So, in conjunction with principles governing rationally holding true, (P5) can figure in explaining why one can't decide to do things that one takes to be mutually incompatible.

Similarly, in conjunction with other principles governing rationally holding true, (P5) can figure in deriving an analogue, for deciding to ϕ , of (P2), Kant's principle, according to which the rational setting of ends requires holding that those ends are in principle attainable. Plausibly, (P7) governs rationally holding things true.

(P7) One can rationally hold that p , and hold that its being true that p entails its being true that q , only if one holds that q .

According to (P7), if one holds that its being true that q is a necessary condition on one's ϕ -ing, then one can rationally hold that one will ϕ only if one holds that q . But we have from (P5) that one can rationally decide to ϕ only if one rationally holds that one will ϕ . So, in the same circumstances, one can rationally decide to ϕ only if one rationally holds that q . Thus, we can derive the following analogue of (P2) for deciding to ϕ :

(P8) One can rationally decide to ϕ only if, with respect to all agential and environmental conditions that one holds to be necessary for its being true that one will ϕ , one rationally holds true that those conditions obtain.

Suppose, for example, that one held that in order for it to be true that one will reach the shops by noon, it would have to be true that one leaves home before 11am. And suppose that one now knew that it was well past 11am. In that case, according to (P8), what one held true would preclude one from now rationally deciding to reach the shops by noon. That is, reason would permit one to decide

to reach the shops by noon only on condition that there were a change in what one knows or holds true about one's circumstances.

The considerations to this point are focused on deciding to ϕ and therefore don't speak directly to (P1) as a principle governing the setting of ends more generally. However, analogous considerations apply with respect to (P1), on the plausible assumption that setting an end is at least approximately equivalent to deciding to try. In both cases, there is a need to explain why certain patterns of decision or end-setting would not be reasonable, and explaining that requires connecting decision and end-setting with positive commitments about what one will do or try to do, and so with what it is possible for one to do or to try to do. A mere absence of commitment to impossibility would not support the required explanation. However, although analogous considerations apply to both deciding and setting an end, there are also important differences between the two cases. In particular, (P1) is weaker than (P5) in a way that means that it does not rule out the possibility of rationally setting incompatible ends. For according to (P1), one can rationally set an end that one holds may not be attained. Since it is possible rationally to hold that it is possible that p and possible that q , even though it would not be rational to hold that it is possible that (p and q), the principle therefore fails to rule out the rational setting of incompatible ends. (See e.g. Bratman 2009.) Now that result is not itself dramatically implausible, although it does mean that one would have to put in additional labour in order to mount a full dress defence of (P1). However, it is worth observing that if we wished to defend the claim that perfect rationality is not consistent with the setting of incompatible ends, then we could impose it as a requirement, without violence to (P1), by the addition of the following plausible seeming principle of unification:

(P9) One can rationally set oneself a totality of ends, Σ , only if one could rationally set oneself as an end the conjunction of all ends in Σ .

The explanatory power of (P1) and (P5) in underwriting the various rational requirements that apply to deciding to ϕ and to setting oneself an end gives us *pro tanto* reasons to endorse those premises. However, those reasons might be defeated by reasons to hold that ordinary thought allows for the possibility of rationally deciding to ϕ whilst failing to hold true that one will ϕ . And they might also be defeated by our failure to make sense of rationally holding true that one will ϕ on the basis of a decision to ϕ . I won't here consider the first potential source of defeaters, but it will be important to say something about the second. (For discussion of the first sort of defeaters, see Holton 2009.) In order to do so, I will outline an account of the nature of deciding to ϕ that has the resources to make transparent the required connection between deciding to ϕ and holding true that one will ϕ . The account to which I'll appeal is one developed and defended by Matthew Soteriou (2013: 257–307). As we'll see, it has some important affinities with some of Kant's thoughts about the topic.

Soteriou's account of deciding to ϕ builds on earlier discussions of ways in which one's deciding now to ϕ can figure in one's deliberation and action, especially discussions by Michael Bratman (1987; 1999; 2006) and David Velleman

(1989; 2000). Crucially, when one has decided to ϕ , one's further practical deliberation will take place under the constraint that one will ϕ . Thus, deciding to ϕ is a way, for practical purposes, of determining an answer to the question, what will one do? For example, in deliberating about what else to do, one will not, without revising or forgetting one's decision to ϕ , consider deciding to do things incompatible with one's ϕ -ing.

For reasons we have already discussed, there are good reasons to think that one's deliberating in that way on the basis of one's having decided to ϕ is to be explained by appeal to one's holding true, on the basis of one's decision, that one will ϕ . That would be consistent with the view, defended by Velleman, that deciding to ϕ is a matter of forming the belief that one will ϕ . However, as Soteriou points out, there are good reasons to think that, unlike paradigmatic cases of belief, where a subject has decided to ϕ ,

...the constraint of treating as true the proposition that she will ϕ is a constraint on the subject's planning that the subject regards as *self-imposed*. (287)

Soteriou, following Velleman (2000: 32–55), explains the way in which a subject will regard their decision as a matter of their imposing constraints on their own deliberation by appeal to the seeming absence of any external source for those constraints:

Acceptance of the truth of the relevant proposition is not one that the subject takes to be grounded in evidence that she possesses, in so far as she takes herself to be epistemically entitled, given her evidence, to make alternative decisions and hence to make any one of a number of other inconsistent assumptions about what she is going to do. So when a subject decides to ϕ , and then subsequently plans on the assumption that she is going to ϕ , the subject assumes something about her own future on the basis of evidence that, *from the subject's own point of view*, simultaneously licenses her to assume something about her future that contradicts it. Note that even *after* she decides to ϕ the subject still takes herself to be *epistemically* entitled, given her evidence, to make an alternative decision, and thereby assume something else about her future. (Soteriou, 2013: 287)

To the extent that the constraints imposed on one's practical deliberation by deciding to ϕ are self-imposed, they are similar to the constraints one imposes on one's theoretical reasoning by supposing something for the sake of argument. Soteriou exploits the comparison in order to illuminate his proposal about the nature of deciding to ϕ :

When you assume that p for the sake of argument, you treat p as true, you regard and treat this constraint on your reasoning as self-imposed, and part of what is involved in treating the constraint as self-imposed is your treating the assumption as one that is to be discharged—e.g. with an outright

conditional judgement that is outside the scope of the supposition. Likewise, when, having decided to ϕ , you plan on the assumption that you will ϕ , you regard and treat this constraint on your planning as self-imposed. And likewise, I want to suggest, part of what is involved in treating the constraint as self-imposed is your treating the assumption as one that it to be discharged. However, in the case of your planning assumption you don't treat the assumption as one that is to be discharged by an outright conditional judgement that is outside the scope of the assumption. Rather, you treat the assumption as one that is to be discharged by the performance of an action that makes the assumption true. (288)

According to Soteriou, then, deciding to ϕ is a matter of imposing on one's own future practical deliberation and activity a specific constraint. In particular, it is a matter of imposing on one's deliberation and activity the constraint of holding that one will make it true that one ϕ s and, so, that one will ϕ . As we've seen, Soteriou holds that, in treating the constraint as self-imposed, we view our holding true as constraining us only via our sustained willingness to remain so constrained. Thus, in deciding to ϕ one is constrained by one's own willingness, rather than by one's evidence. Soteriou's account of the power to decide thereby fits Kant's own conception of an autonomous will. For according to Kant, the autonomy of a person's will amounts to its being

a free will which, in accordance with its universal laws, must necessarily be able at the same time *to agree* to that to which it is to *subject* itself. (1788: 5: 132)

On this view, then, deciding to ϕ is a matter of imposing on one's own practical deliberation the constraint that one will ϕ . Similarly, but more generally, we can treat setting an end as a matter of imposing on one's own practical deliberation the constraint that one will strive to attain that end. Thus, we have presented an intelligible account of why, and how, principle (P1), repeated here, can be true.

(P1) One can rationally set an end, E , only if one rationally holds true that E is in principle attainable.

5. *Faith and evidence.* Although Soteriou views our decisions to ϕ as self-imposed and, so, to that extent as unconstrained by evidence, his account of the way in which one is to discharge one's self-imposed planning assumption serves to bring our decisions into contact with our activities and, thence, with our evidence:

If you take yourself to be incapable of performing the action, then you take yourself to be incapable of making an assumption that is to be discharged in this way. That is why in the case of one's planning assumptions, the range of assumptions one can make is rationally constrained by one's evidence—in particular, one's evidence about what one can and cannot do. (288)

Soteriou's elegant account of deciding to ϕ thus has the resources to make transparent the connection characterized in (P5) between deciding to ϕ and holding true that one will ϕ . The picture is one in which, prior to deciding what to do, one is faced with an array of evidence about one's capabilities in the circumstances. One is faced with evidence to the effect that there are some things one cannot do, and so some things that one will not do. If cognizant of that evidence, one is rationally foreclosed from deciding to do those things. One is also faced with evidence to the effect that there are some things one can do. Typically, one is so presented with a range of things that one can do, and one's evidence fails to dictate which amongst them are things that one will do. Which, if any, of those things one will do is, to that extent, up to oneself. One therefore takes oneself to be free, with respect to any ϕ in that range, to decide whether or not to ϕ . And although one treats a decision to ϕ as constraining one's future practical deliberation, one views that constraint as one that is self-imposed, and so as one that one will respect only to the extent that one is willing to adhere to it. So one does not treat such a decision to ϕ as affecting one's evidential situation, by placing one's not- ϕ -ing outside the range of actions that one is then in a position to undertake. Thus, one would not treat one's decision to ϕ as underwriting an ordinary theoretical belief that one will ϕ .

There are things that one has evidence that one cannot do. One cannot decide to do those things. There are things that one has evidence that one can do. One can decide to do those. To a good first approximation, reasons for belief suffice for reasons to hold things true. Given our epistemic limitations, there are liable also to be things about which one lacks evidence either way: potential courses of action for which one lacks evidence that one cannot complete them and also lack evidence that one can complete them. Where one's ϕ -ing would fall within that penumbral region, is it possible for one rationally to decide to ϕ , where that would be a matter of one's holding true, by way of a planning assumption, that one will ϕ ? In effect, the question concerns whether, in addition to reasons for belief being sufficient for holding true, they are also necessary.

The way we earlier approached this type of question was via the following line of thought. What determines what one will do is a combination of what one has the power to do, given one's circumstances, and how one will make use of that power. What one has the power to do, given one's circumstances, is not up to one. But it is up to one how one will make use of that power. Forming a reasonable view of what one will do, therefore, requires forming a reasonable view both of one's power to act, given one's circumstances and of how one will make use of that power. In making a decision to ϕ , and so in imposing a constraint on one's practical deliberations, one aims to form a reasonable view about what one will do that is based upon a combination of, on one side, information about one's power to act, given one's circumstances and, on the other side, one's commitment to act in a particular way, given that information, in order to discharge the assumption that constitutes one's view about what one will do. One's view now about what one will do in the future can therefore turn out to be wrong on either of two grounds: one might have been wrong about what one had the power to do, given one's circumstances; and one might lose one's commitment to act, by

ceasing to impose on oneself the constraint of seeking to discharge one's initial assumption about how one will act, either through changing one's minds, or through failing to remember what one had decided to do. That is, one might be wrong about either environmental or agential conditions.

Now if we hold that the extent of the space within which one is entitled freely to impose constraints on one's own practical deliberation is determined by what one is in a position freely to make true, then we will hold that that space is determined, in turn, by facts about one's power, given one's circumstances. And one will hold, moreover, that, since one is not in a position freely to determine facts about one's power to act, one is not in a position freely to impose constraints on one's own practical deliberation that amount to holding true specific propositions about that power. Rather, insofar as one's decision to ϕ involves one's holding things true about one's power to act, one's holding those things true will have the same status as one's other beliefs about things outside one's active control, and will therefore be subject to whatever evidential demands such beliefs are subject. (That was the function of (P3) in the argument that was sketched in section 3.) But it is plausible that ordinary beliefs about things outside one's active control are reasonable only if one possesses positive evidence in favour of those beliefs. That is, it is plausible that one cannot reasonably believe that p when one merely lacks evidence in favour of believing that it's not the case that p . Rather, it is plausible that one should instead withhold belief. (That was the function of (P4) in the argument sketched in section 3.) Thus, in particular, in cases in which one's decision to ϕ involves one's holding things true about the prospective activities of other agents, one will require evidence that those activities will be forthcoming—evidence concerning the willingness and competence of those other agents. In cases in which one lacks such evidence, one should withhold belief, and so should not hold true, that the required activities of others will be forthcoming.

We are now in a position to see what is wrong with that line of reflection. It presupposes that holding things true for practical purposes is to be modeled on holding things true for theoretical purposes. Thus, it presupposes that holding things true for practical purposes is subject, given its aims and functions in our cognitive economy, to all of the same demands as paradigmatic—theoretical—belief. However, we have already considered ways in which the epistemological perspective that we have on what we hold true for practical purposes is quite different from the perspective that we have on what we hold true for theoretical purposes. Our holding things true for practical purposes differs from our holding things true for theoretical purposes with respect both to what entitles us to hold things true and to that which our holding things true entitles us. Thus, we are entitled to hold things true for practical purposes on the basis of what we regard as being a free decision, so a decision we regard as one that, consistently with our evidence, we might have failed to make. Moreover, even having made a decision, and so having come to hold something true for practical purposes, we do not regard ourselves as having thereby provided grounds for a theoretical belief about what we will do. For as we've seen, our theoretical beliefs plausibly provide constraints on what we can then decide to do, so that if our theoretical beliefs are inconsistent with its being true that we will ϕ , then we cannot, without changing

our theoretical beliefs, decide to ϕ . So, if deciding to ϕ required forming a theoretical belief to the effect that one will ϕ , then, having once decided to ϕ , we would be incapable of revoking that decision, without giving up a belief that is supported by evidence. So, insofar as one has a conception of belief as governed by evidence, one has reason not to conflate holding things true for practical purposes with holding things true for theoretical purposes. Thus, at least one of premises (P3) and (P4) in the argument sketched in section 3 should be rejected. (As noted earlier, it might be proposed that we should expand the operative notion of evidence so that freely deciding to ϕ can be a way of freely changing one's evidence. But that would serve to support the problematic evidential requirement on what one rationally can hold true for practical purposes only in conjunction with supplementary premises. In particular, it would need to be shown that one's capacity freely to change one's present evidence is itself confined to decision-independent evidence.)

Kant gestures towards a distinction between holding things true for practical purposes and holding things true for theoretical—or speculative—purposes in the following passage:

Thus, in the union of pure speculative with pure practical reason in one cognition, the latter has primacy.... For, without this subordination a conflict of reason with itself would arise, since if they were merely juxtaposed (coordinate), the first would of itself close its boundaries strictly and admit nothing from the latter into its domain, while the latter would extend its boundaries over everything and, when its need required, would try to include the former within it. (1788: 5:121)

Kant's thought here is the following. Suppose that holding things true for practical purposes and holding things true for theoretical purposes constituted a single cognitive kind. In that case, we would be liable to be presented with candidates for being held true with respect to which practical and theoretical reason conflict. For there would be liable to be candidates for being held true such that theoretical reason requires us not to hold them true—since we lack appropriate evidence for holding them true—whilst practical reason requires us to hold them true—since its aims require setting an end that depends upon holding them true. Kant's proposal for avoiding otherwise irresolvable conflicts of that sort is to distinguish holding things true for practical purposes from holding things true for theoretical purposes. In the context of defending specifically moral faith, Kant draws the distinction in the following way:

What belongs to duty here is only the striving to produce and promote the highest good in the world, the possibility of which can therefore be postulated, while our reason finds this thinkable only on the presupposition of a supreme intelligence; to assume the existence of this supreme intelligence is thus connected with the consciousness of our duty, although this assumption itself belongs to theoretical reason; with respect to theoretical reason alone, as a ground of explanation, it can be called a *hypothesis*; but in relation to the intelligibility of an object given us by the

moral law (the highest good), and consequently of a need for practical purposes, it can be called *faith* and, indeed, a pure *rational faith* since pure reason alone (in its theoretical as well as in its practical use) is the source from which it springs. (1788: 5: 126)

So Kant proposes that what we hold true for practical purposes is to be viewed, from a purely theoretical perspective, as a mere hypothesis. That allows that such a holding true can be detached from the requirement for positive evidential support whilst, at the same time, leaving open that rationally holding something true would not be possible in the face of positive evidence against doing so. For one can be entitled rationally to accept hypotheses in the absence of positive evidence that they are true; but one must be prepared to reject them in the face of positive evidence that they are not true. And we've considered reasons for thinking that deciding to ϕ , and so for holding true for practical purposes that one will ϕ , are subject to an analogous pattern of requirements. (Kant discusses in more detail the government of rational hypotheses in 1781/1787: A769/B797–A782/B810. For useful discussions of Kant's views about the nature of holding true in general, as well as for practical purposes, see Chignell 2007 and Stevenson 2003.)

Our ordinary view is that it can be perfectly rational to decide to ϕ in the absence of evidence that we will ϕ , and thus that we can set ourselves an end in the absence of evidence that the end is in principle attainable. Thus, our ordinary view fails to support the combination of premises (P3) and (P4) in the argument sketched in section 3. I've drawn on the work of Soteriou and Kant in order to present a theoretical perspective on decision—and, so, on holdings things true for practical purposes—that is able to sustain our ordinary view. When one decides to ϕ , one thereby holds true that one will ϕ . If one realizes that one will ϕ only if certain environmental conditions obtain, then one's rationally holding that one will ϕ requires that one hold that those environmental conditions do obtain. Similarly, when one decides on an end, one thereby holds true that one will strive to attain that end, and that the end is in principle attainable. If one realizes that an end is attainable in principle only if certain environmental conditions obtain, then one's rationally holding that the end is attainable in principle requires that one hold that those environmental conditions do obtain. Thus, in particular, in cases in which one knows that the attainability of one's end is dependent on the activities of other agents, one must hold true that those activities will be forthcoming. However, holding things true for practical purposes differs in nature from holding things true for theoretical purposes, and so is subject to different requirements. The differences make space for rationally holding things true for practical purposes in cases in which couldn't rationally hold them true for theoretical purposes because one lacked suitable evidence. So, the differences makes space for the possibility of rationally holding true that the activities of other agents on which one's successful attainment of an end depend will be forthcoming in the absence of evidence that those activities will be forthcoming.

At the end of section 3, I suggested that a full defence of our ordinary view would require a response to the claim that holding true is governed by the constitutive aim of holding true only what is true. We are now in a position to see

that such a response would require an alternative account of the constitutive aims of the setting of ends and, so, of holding things true for practical purposes. However, our discussion has, from the outset, been shaped by just such an account. For Kant's account of practical reason is one according to which its constitutive aim is a moral aim: the attainment of the highest good.

Our ultimate moral end, the highest good, imposes upon us the need to set ourselves mediate ends, what we take to be means to our ultimate end. Our practical and theoretical limitations mean that the mediate ends that we set are liable to instability. For instance, our conception of our ultimate end is limited so that it leaves open various questions concerning the space between what is morally obligatory for us and what is morally impermissible for us. It leaves open some questions about the precise boundaries of the space. And, in addition, it leaves open questions about our optimal route through it. Thus, with respect to any mediate end that we set, it may turn out that between setting the end and attaining it, we will cease to constrain ourselves by holding true that we will strive for that end. Moreover, limited as we are, we may simply forget some amongst the ends that we have set. So, from a broadly practical perspective, our ends are liable to a certain amount of instability.

Our epistemic limitations mean that our mediate ends are liable to an additional layer of instability. For we may often need to set mediate ends that we do not know to be attainable. With respect to any of those mediate ends, it may turn out that between setting the end and attaining it, we will acquire evidence that it is in principle unattainable. If we were to acquire such evidence, then we would be required to shed our adherence to the constraint of holding true that we will strive for the end. If we set only those ends that we knew to be in principle attainable, our striving towards those ends would be to that extent more stable, and we would be more likely to attain those ends. But the gain in stability and likelihood of success would have to be traded off against the limitations our adhering to that restriction would impose upon the space of our options. If our most fundamental aim in setting ends were the attainment of those ends, then the gain in likelihood of success would be privileged. That might be so, for example, if our fundamental aim in setting ends were the attainment of knowledge about what we will do or will strive to do. However, insofar as our ultimate practical aims impose distinctive requirements on what we should strive to attain, and so force us to set ends outside the space of what we know to be attainable, those aims are liable to outweigh our concern to obtain a prior guarantee of stability and success.

If one knew now that a particular outcome will be forthcoming in the future, one would be in a position confidently to rely on that outcome. In particular, one would be in a position to be confident that one will not later come across decisive evidence that the outcome will not occur. More generally, one would be in a position to be confident that one's current view about the future need not be dislodged by genuine reasons. One might lose one's view in other ways—for example, by losing confidence in the face of merely apparent counter-evidence, or simply by forgetting—but one will not be forced to shed it by reason alone. For reasons we've just considered, one typically lacks any similar guarantee with respect to what one decides to do or to strive for. But we can now see that a

closely similar guarantee is available, according to Kant, with respect to our setting as our end the attainment of the highest good. For according to Kant, neither of the two potential sources of instability in our ends applies there. The practical source does not apply, since there cannot be genuine practical reasons for shedding the end of attaining the highest good. And the theoretical source does not apply, since there cannot be genuine evidence that the highest good is unattainable. Thus, one can be as confident that one's current view about the attainability of the highest good need not be dislodged by genuine reasons as one could be if one knew—*per impossibile*—that the highest good were attainable. Holding true for practical purposes that the highest good is in principle attainable is, therefore, distinctively stable and secure. However, that does not indicate any unreasonableness in more ordinary cases of holding things true for practical purposes. Rather, it indicates only that the reasonableness of our holding things true for practical purposes in more ordinary cases is dependent on our being appropriately sensitive to the potential instability and insecurity of our so doing.

Let's return briefly to the three questions raised in section 2:

Q1. What cognitive requirements must one meet if one is rationally to set an end? Must one, for example, hold true that the end is in principle attainable?

Q2. What is the nature of the attitude of holding true that figures in meeting those cognitive requirements? Is it a paradigmatic form of belief, or does it take some other form?

Q3. To what extent can the required form of holding true be sustained rationally by the needs of practical reason? Does the answer to that question depend on whether the operative needs arise from an absolute demand on practical reason, that we follow the moral law? And does it depend on whether the required form of holding true concerns only things about which we cannot know?

My answer to Q1 has been that rationally setting an end requires holding true that the end is in principle attainable. That requirement is favoured over the weaker requirement that rationally setting an end requires failing to hold true that the end is not in principle attainable by its superior capacity to explain constraints on the rational combination of the setting of ends. My answer to Q2 has been that the required attitude of holding true is not a paradigmatic form of belief. And my answer to Q3 has been that such holding true is grounded practically rather than evidentially. Typically, such holding true potentially is subject to practical or evidential defeat, and so lacks the stability of knowledge. Kant's more exigent demands on faith—according to which it is sustained by absolute moral demands and is insusceptible to evidential disconfirmation—are designed to align faith with knowledge; there is no reason to expect that alignment to obtain with respect to quotidian forms of faith or trust.

6. *Conclusion: faith and trust.* According to Kant, we can decide to do things, and thus impose on ourselves the constraint of holding true that we will do those things, on practical grounds. And we can decide to strive to attain things, and thus impose on ourselves the constraint of holding true that those things are in principle attainable, on practical grounds. My aim to this point has been to make space for Kant's otherwise plausible view by divorcing that view from some dubious near relatives. Thus, I've explained that such a view needn't be committed—and Kant's view is not committed—to the claim that one can decide rationally to do things, or strive for things, that one knows are unattainable. And I've explained that such a view needn't be committed—and Kant's view is not committed—to the claim that in deciding to do things or strive for things on practical grounds one is thereby coming to believe on practical grounds that one's ends are attainable. In doing so, I appealed to Kant's specific account of the constitutive aim of holding things true for practical purposes, according to which that aim is the attainment of the highest good. However, the same outcome would be achieved if it could be shown, more minimally, that the constitutive aim of holding things true for practical purposes is distinct from the constitutive aim of belief. (Independently, there need be no commitment to the claim that one can decide to believe things, for there need be no commitment to thinking that, simply because deciding to ϕ entails holding true that one will ϕ , it follows that deciding to ϕ entails *deciding* to hold true that one will ϕ . So, even if deciding to ϕ on practical grounds entailed believing that one will ϕ on practical grounds, it need not entail in addition *deciding* to believe that one will ϕ on practical grounds.)

With those resources in hand, let's return to the question about trust with which we began. Can it ever be reasonable to trust other people, or to rely upon them, without evidence of their trustworthiness or reliability? The answer we have to this point is conditional: yes, provided, first, that there is not too much evidence that they are untrustworthy or unreliable and, second, that there are genuine practical reasons for trusting them or relying upon them. However, it seems obvious that we are constantly presented with genuinely practical reasons for trusting others. That is, it seems obvious that we are constantly presented with valuable ends the attainability of which is dependent on others. And typically, at least, our situation is one of lacking evidence for others' trustworthiness or reliability, rather than possessing evidence for their untrustworthiness or unreliability. Plausibly, then, it can sometimes be reasonable to trust other people, or to rely upon them, in the absence evidence of their trustworthiness or reliability. Of course, that leaves open the extent to which trust is sometimes made reasonable in other ways—for example, by the obtaining evidential grounds for holding true (in that case, plausibly, for believing) that someone is reliable or trustworthy. That is, it leaves open that a full account of the bases of trust might need to include a response of the third type considered in section 1, an attempt to argue that our epistemic limitations are less stark than the puzzle about trust makes them out to be. And it leaves open that there may be cases in which the potential cognitive consequences of trusting someone are dependent on the obtaining of such evidential grounds—for example, cases in which those consequences include knowledge or belief and, so, cannot straightforwardly be sustained by practical grounds for holding things true. However, my aims here

have been limited to sketching out some space for the fourth response to the puzzle, by indicating ways in which trust might sometimes be reasonable in the absence of evidential grounds for belief. I will conclude by raising two more delicate questions about the reasonableness of trust and reliance, questions that concern the balancing of evidential grounds for distrusting others against practical grounds for trusting them.

It seems obvious that we must often try to balance such grounds. Kant himself expressed some pessimism about the results of the trade off:

We must so conduct ourselves to a friend, that it does us no harm if he were to become our enemy; we must give him nothing to use against us. We are not, indeed, to suppose that he may become our enemy, for then there would be no trust between us. But if we give ourselves entirely to a friend, and entrust him with all the secrets which might detract from our happiness, and might well be divulged if he did become an enemy, then it is very unwise to tell him these things, since he could either give them away through inadvertence, or use them to our hurt if he became our foe. (Kant 1784–5: 27: 429–430)

The two questions that I wish to table concern trust in other people's good disposition of will, rather than in other competences of theirs on which we might rely. The first question is this. Could there be cases in which trusting someone is made entirely unreasonable on evidential grounds? That is, could we know enough about someone's poor character that no countermanding practical grounds, including grounds deriving from morality, would outweigh what we know and thus make it reasonable for us to trust them? The answer to this question will depend ultimately on answers to further questions, concerning the extent to which others may be regarded as free, not only from their own practical perspective, but also from our practical perspective. Can we have sufficient evidence about the dispositions of others' wills to rule out their striving to do what they are trusted to do? Could such evidence ever foreclose on their absolute freedom to go against all past evidence concerning their character and thus redeem themselves by deciding to warrant our trust?

Second, could there be cases in which practical grounds would make it obligatory to trust someone, at least with respect to the disposition of their will? Despite Kant's animadversions about apportioning complete trust in a friend, given the then current state of human relations, we might reasonably wonder whether there might nonetheless be moral grounds for striving to trust others, insofar as doing so is made possible by our evidence. And we might wonder, in particular, whether the highest good could comprise a realm in which every person is in fact fully trustworthy with respect to the disposition of their wills, since their wills are disposed in accord with respect for the moral law, and yet in which not every person is fully trusting. Thus, we might wonder whether fully respecting the demands of morality, and so setting the highest good as one's end, in turn imposes on one a requirement to strive to have faith in other people, insofar as one isn't precluded from doing so by evidence of their liability to malevolence.

References

- Bratman, M. 1987. *Intentions, Plans, and Practical Reasoning*. Cambridge, MA: Harvard University Press.
- . 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.
- . 2006. *Structures of Agency: Essays*. Oxford: Oxford University Press.
- . 2009. “Intention, Belief, Practical, Theoretical.” In S. Robertson (ed.) *Spheres of Reason*. Oxford: Oxford University Press.
- Chignell, A. 2007. “Belief in Kant.” *Philosophical Review* 116, 3: 323–360.
- Gardner, S. 2006. “The Primacy of Practical Reason.” In G. Bird (ed.) *A Companion to Kant*. Oxford: Blackwell.
- Holton, R. 2009. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Kant, I. 1781/1787. *Critique of Pure Reason*. P. Guyer and A. W. Wood (ed. and trans.). Cambridge: Cambridge University Press, 1998.
- . 1784–5. “Moral Philosophy: Collins’s lecture notes.” In *Immanuel Kant: Lectures on Ethics*. P. Heath (ed. and trans.) and J. B. Schneewind (ed.). Cambridge: Cambridge University Press, 1997.
- . 1786. “What does it mean to orient oneself in thinking?” In *Religion within the Boundaries of Mere Reason*. A. Wood and G. di Giovanni (ed. and trans.). Cambridge: Cambridge University Press, 1998.
- . 1788. *Critique of Practical Reason*. M. Gregor (ed. and trans.). Cambridge: Cambridge University Press, 1997.
- . 1790. *Critique of the Power of Judgment*. Paul Guyer (ed. and trans.) and E. Matthews (trans.). Cambridge: Cambridge University Press, 2000.
- . 1793. *Religion within the Boundaries of Mere Reason*. A. Wood and G. di Giovanni (ed. and trans.). Cambridge: Cambridge University Press, 1998.
- Kleingeld, P. 1998. “Kant on the Unity of Theoretical and Practical Reason.” *The Review of Metaphysics* 52, 2: 311–339.
- Neiman, S. 1994. *The Unity of Reason: Rereading Kant*. Oxford: Oxford University Press.
- Soteriou, M. 2013. *The Mind’s Construction: The Ontology of Mind & Mental Action*. Oxford: Oxford University Press.
- Stevenson, L. 2003. “Opinion, Belief or Faith, and Knowledge.” *Kantian Review* 7: 72–101.
- Timmermann, J. 2009. “The Unity of Reason: Kantian Perspectives.” In S. Robertson (ed.) *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Velleman, D. 1989a. *Practical Reflection*. Princeton, NJ: Princeton University Press.
- . 2000. *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- Watkins, E. 2010. “The Antinomy of Practical Reason: reason, the unconditioned and the highest good.” In A. Reath and Jens Timmermann (eds.) *Kant’s Critique of Practical Reason: A Critical Guide*. Cambridge: Cambridge University Press.

- Willaschek, M. 2010. "The primacy of practical reason and the idea of a practical postulate." In A. Reath and Jens Timmermann (eds.) *Kant's Critique of Practical Reason: A Critical Guide*. Cambridge: Cambridge University Press.
- Wood, A. W. 1970. *Kant's Moral Religion*. Ithaca, NY: Cornell University Press.