

teorema

Vol. XXXVII/2, 2018, pp. 165-180

ISSN 0210-1602

[BIBLID 0210-1602 (2018) 37:2; pp. 165-180]

The Influence View on Responsible Belief Reply to Kulp, Levy, Rossi, and Goldberg

Rik Peels

I. INTRODUCTION

Let me start by thanking all the authors for their meticulous, clear, and rigorous criticisms. I immensely enjoyed reading their papers and I am deeply honored by the fact that these philosophers have spent so much time and energy on what, no doubt, would have cost them much less time and energy if only my points had been clearer and better argued. I have also learned a lot from these authors. Their essays made me rethink especially four things: (1) to what extent beliefs can be voluntary in relationships in which the truth of a proposition depends on whether or not it is believed, (2) how to define compatibilism in the free will debate and in the ethics of belief debate and the differences between these two, (3) how to differentiate between beliefs that we have influence on but for which we *are* responsible on the one hand, and beliefs on which we have influence but for which we are *not* responsible (if there are such beliefs) on the other, and (4) the exact relationship between epistemically responsible belief on the one hand and epistemically justified belief and knowledge on the other. I hope to delve into these issues in more detail on future occasions.

That being said, I'm not entirely convinced by the main criticisms. Below, I address the core worries of each of the authors. I have not been able to address everything, but I hope that tackling the most important issues will help move the debate forward.

II. BELIEVING AT WILL: A REPLY TO KULP

In his contribution, Christopher Kulp provides a myriad of fascinating cases, both of responsible belief and of blameworthy belief. In the

final section of the paper, he presents four principles about belief formation. Let me say straight away that I agree with each of those and that I find them helpful in construing and assessing an ethics of belief. Here, I focus on where we disagree and on where Kulp takes us to disagree. In doing so, I confine myself to *propositional* belief even though I agree with Kulp it would be interesting to apply my Influence View of responsible belief to *objectual* belief (I leave doing so for another occasion).

In chapter 2 of *Responsible Belief* [Peels (2017)], I argue that we should not explain doxastic responsibility in terms of intentional control over our beliefs. That is because, with a few exceptions, we cannot form an intention to form a belief and then carry out that intention — this in opposition to, say, actions: we *can* form the intention to perform some action *A* and then carry it out. With many others, I call this thesis ‘doxastic involuntarism’, a view that has also been defended by William Alston (1989). According to Kulp, my examples in favor of doxastic involuntarism are rather one-sided in that I focus on belief with respect to propositions *that are obviously true*, such as that $2+2=4$ and that this is my hand in front of my face as I write this. I do not think this is entirely right: in §2.4 of the book, pp. 61-64 [“the book”, unless indicated otherwise, refers throughout the paper to Peels (2017)], I discuss in detail an influential example given by Matthias Steup (2017) about a parked car in which my evidence with respect to a specific proposition is on balance. I agree, though, that it would be helpful to discuss more examples in which one’s evidence does not clearly count in favor of belief or in favor of disbelief.

Now, Kulp rightly points out that there is an important class of cases that I do *not* discuss in the book, a class famously introduced by William James (1897) p. 57, Kulp asks us to imagine that Jones is involved in a mountain climbing accident and that Jones has good reason to think that climbers’ prospects of survival are greatly enhanced by having a positive attitude about their capacity to survive. He argues that in such a case (he adds some further conditions), Jones can voluntarily or intentionally believe or not believe the following two propositions:

- (1) I shall survive this dire mountain climbing emergency, *even though my evidence is that there is a slightly less than even chance that I will make it.*
- (2) I shall survive this dire mountain climbing emergency, *even though my evidence is that there is only a 50/50 chance that I can make it.*

Now, my reply to this is rather brief: I think Kulp is *exactly right*. This is why on p. 55 of the book, I endorse doxastic involuntarism only as a contingent, psychological thesis that is largely true, but *not* as a conceptual thesis. By that, I mean that it is not a conceptual truth that we cannot believe at will. It is merely a psychological truth that applies to most cases. In fact, I have spelled out Jamesian cases of believing at will in much more detail elsewhere and argued against various objections to them [see Peels (2015)]. Thus, in certain cases in which truth depends on belief, it is up to one whether or not one believes the relevant proposition.

Kulp gives some further examples of truth-depends-on-belief [TDB-cases, as I call them in Peels (2015)], e.g. the belief that someone else is a very good pitcher but that one can hit his fastball. In the book, I pay no attention to such cases (I now think I should have made explicit why I do not do so). The reason for that is that they *cannot explain widespread doxastic responsibility* and that is because truth-depends-on-belief cases are relatively rare. In most cases in which people seem responsible for their beliefs, the truth of their belief is independent from whether or not they believe it.

Kulp's insightful essay made me rethink, though, how rare such truth-depends-on-belief cases really are. Upon reflection, it seemed to me that there is a rather large class of cases that neither he nor I take into account, namely all sorts of beliefs *in personal relationships*. There are quite a few cases in which the relationship will work only if some proposition for which the evidence is not convincing is believed, for example, such propositions as: the other person loves me (in certain romantic relationships), the other person is trustworthy (in virtually every relationship), this friendship will be life-long, and so on. I think Kulp has rightly drawn our attention to this; it deserves much more philosophical consideration.

Now, Kulp claims that there are further counter-examples to the thesis of doxastic involuntarism. First, he sketches a scenario in which Smith, a mountain guide, in a dangerous scenario *chooses* to believe the following proposition:

- (3) My evidence is 50/50 that the snow slope will avalanche, but it is safe to cross.

Kulp takes it that Smith's belief is irresponsible — or, more precisely, *epistemically* irresponsible. After all, unlike Jones, Smith has no reason to think that his belief itself will contribute to the likelihood that the belief is true. His point is that my account of responsible belief cannot make sense of the fact that Smith's belief is irresponsible, since I deny doxastic

voluntarism. I'm not convinced this should count as a case of voluntary belief, though. There is no reason to think that in this case, Smith truly chooses to *believe* proposition (3). A more plausible interpretation, it seems to me, is that he chooses to *accept* – that is, act and reason from the assumption that – proposition (3) is true. The distinction between belief and acceptance has been famously spelled out by Cohen (1992).

Kulp goes on to sketch three fascinating examples that involve a child's health rather than a mountain climbing accident. Since he believes my account of responsible belief can make sense of cases #1 and #2, but not of case #3, let us focus on case #3, which runs as follows. Black's child suddenly begins to choke. Black knows the Heimlich Maneuver and believes the evidence is 50/50 that she can benefit her child by performing the procedure. She *chooses to believe* that she can successfully perform the Heimlich Maneuver on her child and, subsequently, successfully carries it out. Kulp suggests that her belief may well be responsible. But my account, Kulp suggests, does not provide the resources to assess whether Black's belief is responsible, since it occludes cases of voluntary belief formation.

Again, though, it seems to me that Black's case is best understood in terms of acceptance. Black has insufficient reason to think that performing the Heimlich Maneuver will save her child, but not doing it is equally bad, so she decides to *go for it*, that is, act on the assumption that it will work. Another plausible interpretation would be to say that Black believes that *if* she believes that performing the Heimlich Maneuver will save her child, that makes it rather likely that performing it will save her child — say, because that makes her perform the maneuver better. But then we are back to a truth-depends-on-belief case. In that case, Kulp's fascinating example would be another reason to think that such truth-depends-on-belief cases are not as rare as one might initially think.

III. DOXASTIC COMPATIBILISM AND THE INFLUENCE VIEW: A REPLY TO LEVY

Neil Levy and I disagree on responsible belief. Not because we disagree on how to spell out responsible belief, but because he believes that there is no such thing as responsible belief. In fact, on his view, there is no such thing as responsibility at all, as he has argued in detail elsewhere [see Levy (2011)]. Understandably, therefore, he focuses on what both of us believe to exist: agency and control. He argues for a different conception of agency and control than the one I defend in *Responsible Belief* and champions the view that if his *synchronic* conception of agency in terms of

doxastic control is in trouble, then so is my alternative, *diachronic* conception of agency in terms of doxastic influence. Unfortunately, I have to leave doing justice to each of his criticisms for another occasion. Here, I focus on his main worries.

As to the terminology, I agree with Levy that I should have defined ‘compatibilism’ more carefully when it comes to the debate about free will: not so much in terms of the *ability to act otherwise*, but in terms of *determinism* (even though the ability to act otherwise is the issue relevant for the book: on p. 72 I define ‘compatibilism’ in terms of the inability to act otherwise). We should note, though, that, contrary to what Levy claims, my use of the phrase ‘doxastic compatibilism’ is *not* at all idiosyncratic. In my book, it refers to the thesis that having obligations to believe does not require voluntary or intentional control over our beliefs. This is exactly how the phrase is used in the debate on doxastic compatibilism. See, for instance, Bayer (2015), Booth (2014), Meylan (2013), Nottelmann (2007), Ryan (2003), Steup (2017), and Wagner (2017). In the book, I note explicitly that in the ethics of belief literature, the term ‘compatibilism’ is used *differently* from how it is used in the philosophy of action and the debate on free will.

In any case, Levy’s own conception of agency is that an agent has control over φ -ing just in case her φ -ing is sufficiently sensitive to reasons, recognizing these reasons as considerations in favor of alterations and adjustments, which is, of course, a matter of degree. Now, in my book I give an example that is meant to illustrate that such control does *not* suffice for doxastic responsibility:

Imagine a possible world in which there are creatures who are like us in that their belief-forming mechanisms are largely functioning properly: upon having the experiences and beliefs we have, they roughly form the same beliefs as we do. In one regard, however, they are crucially different from us: they cannot *influence* what they believe. Thus, they cannot gather evidence, work on their intellectual virtues and vices, improve the functioning of their cognitive mechanisms, and so forth. For instance, they cannot decide to think about something or reflect on their reasons to believe something. Would we hold those creatures responsible for their beliefs in this *recherché* scenario? It seems clear to me that we would not. Their belief formation is not up to them. Their beliefs are simply the deliverances of their cognitive mechanisms in combination with certain inputs. But, we have assumed, neither the functioning of their cognitive mechanisms nor the scope or quality of their evidence base is up to them. It seems clear that if these are not up to them, the output is not up to them either and it would be unfair to hold them responsible for their beliefs [Peels (2017), pp. 74-75].

Levy responds that intentions (or decisions) do not add anything here when it comes to control and that I seem to be after controlling our control systems. However, I am not. In fact, the agents in the example clearly have *some kind of control*. I am concerned though with *control that suffices for responsibility and their control*, so I argue, fails to be sufficient for that. In order for them to be responsible for their beliefs, they also need to be able to *influence* their beliefs, by intentionally, say, gathering further evidence. Mere reasons-responsiveness will not do for doxastic responsibility.

Now, he goes on to say that the creatures imagined are not merely the puppets of their doxastic mechanisms. After all, he says, they are *constituted by* these mechanisms. I agree that they are indeed partly *constituted* by these mechanisms, but that does not suffice to escape being a puppet: robots are partly constituted by their cognitive and non-cognitive mechanisms, but they are not truly *agents*. We are partly constituted by our digestive systems and our automatic fear-responses in response to dangerous situations, but that does not suffice to make us agents.

Let us now turn to another worry that Levy addresses. In chapter 4 of *Responsible Belief*, I argue that in most cases, the evidence we have forces us to have certain beliefs. In such cases our will or intentions or decisions do not and cannot make a difference to what we believe. Levy argues that he finds this use of the phrase ‘force’ misleading. In the book, I point out, though, that I use the word ‘force’ somewhat stipulatively [p. 134]. Being forced by the evidence to hold a belief is not like being forced by a kidnapper to remain seated: our own mechanisms play a crucial role in the former, but not so much in the latter case. Still, it seems to me that we can rightly use the word ‘force’ in this case, since there is a relevant relation between one’s evidence and one’s will: if there is nothing one could have done to change one’s evidence and doxastic mechanisms and if, in virtue of one’s evidence and doxastic mechanisms, one inevitably forms a belief, it does *not* seem right to hold that person responsible for holding that belief (chapter 4 of the book is, after all, a chapter on the conditions under which one is or is not excused for holding a belief).

Another worry brought forward by Levy concerns my claim, in response to the so-called Tracing Problem, that people can be held responsible for violating an intellectual obligation despite not acting against any *occurrent* beliefs, because they can also be acting against their *dormant* or *tacit* beliefs (pp. 193-194; see also Goldberg’s contribution in this journal issue). Levy suggests that such dormant and tacit beliefs are, by hypothesis, offline. Subsequently, he wonders how agents can be properly held responsible for not acting on their dormant and tacit beliefs if they

are not properly cued for bringing those beliefs online (he spells out this point in more detail in Levy (2014); (2016)). This is a complex issue that I cannot discuss in any detail here. But let me point out that people *can and sometimes should cue themselves*. By being careful and by paying attention, they cue themselves to bring certain beliefs online. We blame people for being inattentive, which can result in the failure to bring various beliefs online. And that seems entirely right, unless certain special excusing circumstances hold, such as being distracted by shootings in the background, being severely ill, and so on.

Finally, Levy argues that if, as I claim, synchronic force is a problem, then diachronic control will not solve it. He asks us to imagine that Chen grew up with narrow-minded parents but became open-minded in the course of her life. Presumably, she at some point came to believe that she was narrow-minded and that that was a bad thing. But that belief, if I am right, was forced upon her, so she is not responsible for that belief or the ensuing action. This is a version of the so-called tracing problem and I discuss it in much more detail in chapter §5.5 of the book [pp. 185-195; see also Peels (2011)].

My reply is briefly this. At some point, Chen may have come to believe that she was narrow-minded and that that was wrong or, maybe, that she could very well be narrow-minded and that that was wrong, or that she was narrow-minded and that that could very well be wrong, or some such thing. But subsequently, given her free will, it was *up to her* whether or not to try to actually *do* something about it or not. Apparently, she met her obligation and became more open-minded, but, presumably, she could have acted from *akrasia* – or performed a series of actions from *akrasia* – so that she would have remained narrow-minded or would have been significantly more narrow-minded. Thus, the fact that our beliefs are synchronically forced upon us does not remove responsibility for the actions we can perform partly based on those beliefs and for the beliefs issuing from those actions. Beliefs are very often *not* causes: what we do, given our beliefs, is up to us.

IV. THE PROBLEM OF DIFFERENTIAL DERIVATIVE RESPONSIBILITY: A REPLY TO ROSSI

Rossi's focus is on the way that I explain doxastic responsibility in terms of the *influence* we have over our beliefs in virtue of our control over such belief-influencing actions as gathering evidence and working

on our intellectual virtues and vices. Thus, on the Influence View I defend, our responsibility for our beliefs is *derivative* from our *original* responsibility for our actions and omissions.

Before we consider his criticisms in more detail, let me make one preliminary point. According to Rossi, I claim that derivative responsibility is a distinct *kind* of responsibility. I think this is mistaken. I acknowledge that to be derivatively responsible is to be responsible *in a different way* than to be originally responsible, but I explicitly deny that these are two *distinct kinds* of responsibility: “These are not two *kinds* of responsibility, but two *ways* of being responsible” [p. 117]. Thus, I am perfectly happy to embrace Rossi’s suggestion in his first footnote which says that the conditions for responsibility are disjunctive: one is responsible for φ -ing just in case one has either *control* over φ -ing (original responsibility) or influence on φ -ing (derivative responsibility).

The main and rather serious worry that Rossi raises for my Influence View on doxastic responsibility is that we need some way to distinguish between consequences that are under our influence for which we *are* responsible and those many other consequences under our influence for which we are *not* responsible. I will follow Rossi in calling this the ‘problem of differential derivative responsibility’. Rossi argues that this is a problem both for my account of *responsibility for belief* and for my account of *responsible belief*. Let us consider them in that order.

Rossi selects one of my examples to illustrate the first point. If it is my task to fill oxygen bottles for the ambulance, but culpably fail to do so, so that a man dies, I am at least partially blameworthy for that. But if his wife remarries and her new husband abuses her children, I am *not* blameworthy for that, even though it would not have happened if I had met my duty by filling the oxygen bottles. Things are similar in the doxastic realm: if Julia has an intellectual obligation that she violates at t and, as a result of that continues to believe that p , then, if at some later time t^* she is afflicted with Alzheimer’s disease and believes at t^* that p , she is not blameworthy for believing p at t^* , even though (so we can stipulate) she would not have believed that p at t^* if she had met her intellectual obligations at t .

I think Rossi’s examples convincingly show that I was mistaken in claiming relatively early on in the book that the following two conditions imply that some person S is responsible for her belief [see Peels (2017), p. 143]:

- (i) S believes that p ;
- (ii) there is some belief-influencing action or series of actions A that S could have performed such that if S had performed A , S would not have believed that p .

Note, though, that *after* this initial account, I refine the account further on in the book. I even explicitly mention the case of children, Alzheimer's patients, and other people whom we cannot hold responsible for their beliefs [p. 183]. I, therefore, add as a necessary condition that the subject in question should at least *have some intellectual obligations* [p. 184]. This means that the Alzheimer case no longer counts as a counterexample. However, it leaves the general problem of differential derivative responsibility for belief intact.

My inclination would be to stress what I say on p. 140 of the book: the conditions for being responsible for a belief are rather low, possibly much lower than being responsible for other consequences of one's actions and omissions. As long as one is subject to intellectual obligations and one could have made a difference to the belief in question but did not violate any intellectual obligation relevantly related to that belief, one is responsible and blameless for holding it. As I argue elsewhere in the book [p. 215], adding a foreseeability condition leads to trouble. I do not think this is a problem: there is a clear difference between being derivatively responsible (and blameless) for beliefs that one could not foresee but *that one has actually formed* on the one hand and such a remote and indirect consequence as the following on the other: the wife of the person who becomes widowed as a result of my action marries a new person and that person happens to abuse her children. Thus, the bar for being responsible for a belief is quite low indeed.

The second problem of differential responsibility that Rossi raises concerns *responsible belief* rather than *responsibility for belief*. My account needs to be able to explain why some beliefs over which we have influence are *responsible* whereas others are *blameworthy*. In the book, I phrase the problem in terms of luck. Rossi thinks that is unfortunate, since something is a matter of luck only if it is significant to someone (something Rossi and I agree on). I would say, though, that *any belief* is significant to the person holding it, since believing the truth has epistemic value and by not believing a truth one misses some kind of epistemic value. But for those who disagree, we can phrase the issue entirely in terms of the absence of control rather than in terms of luck.

His first worry here is that my account seems to imply some kind of doxastic *original sin*. Once you have violated an intellectual obligation and as a result of that, say, continue to believe that p while you would have abandoned p if you had met your intellectual obligation, you remain blameworthy for believing that p . This is true no matter what you do later on in your life, such as maintaining your belief that p because a friend whom you deem highly reliable tells you that p is true. He concludes that it is *not* necessary for responsible belief that you have not violated any intellectual obligation in coming to hold or in maintaining that belief. I agree this is a challenging example. Maybe what we should say is that if one performs an action later in life that amounts to something sufficiently similar to meeting the original intellectual obligation, one's belief is blameless — but then it seems unlikely that in that case one would continue to hold that belief. If it is *not* sufficiently similar — as in the above case of a friend's testimony — the evidential basis of one's belief has somewhat improved and that is a laudable thing, but since one still would not have had that belief if one had met one's intellectual obligation, it seems one is blameworthy for having it.

Rossi also argues that it is not *sufficient* for responsible belief that you have not violated any intellectual obligation in coming to hold or in maintaining the belief in question. Imagine, he says, that Jenny wants to become a computer-controlled machine tool operator and that she is deeply ignorant regarding all sorts of propositions as to how to operate the machine, including proposition q . She is under an intellectual obligation to study the relevant material in order to become such a machine operator. She culpably fails to do so and consequently comes to believe that q is false. Thus, she seems blameworthy for believing q is false. On my account, however, Jenny is blameworthy for believing q is false only if q is non-accidentally related to her intellectual obligation.

In reply, let me point out that it seems that q is indeed non-accidentally related to Jenny's intellectual obligation. After all, failing to believe q is objectively bad in that being a computer-controlled machine tool operator comes with knowledge of q and trying to become a tool operator comes with studying q . Rossi seems to mistakenly think that a belief is only non-accidentally related to an intellectual obligation *if one actually holds that belief*. That is false: what matters is whether one has an intellectual obligation partly in virtue of the fact that there is something *subjectively or objectively* bad about not believing q [see p.f 217 of the book]. And given that becoming a tool operator comes with knowledge of q , that is clearly the case.

V. RESPONSIBLE BELIEF, EPISTEMIC JUSTIFICATION, AND
INTELLECTUAL OBLIGATIONS: A REPLY TO GOLDBERG

I thank Sanford Goldberg for his lucid and astute criticisms and I wholeheartedly accept his crisp summaries of my views. I am also glad that he is somewhat sympathetic towards a *diachronic* view of doxastic responsibility in terms of influence rather than the *synchronic* compatibilist view in terms of control that is so popular nowadays (even though his focus is on epistemic justification and not so much on doxastic responsibility). If the diachronic view is right, whether or not one believes responsibly depends on all sorts of belief-influencing actions and omissions that led up to that belief. And I wish I had paid as much attention to the social situatedness of epistemically responsible belief as Goldberg does in his recent book [Goldberg (2018)].

Now, at the outset of his reply Goldberg distinguishes between two projects: (1) analyzing *epistemically responsible belief*, which is somehow relevant for epistemic justification and knowledge, and (2) analyzing *doxastic responsibility*. I agree that there are two distinct projects, but I would characterize them somewhat differently, namely: (1) analyzing *epistemically justified belief*, which is somehow relevant for knowledge, and (2) analyzing *epistemically responsible belief*. I am primarily concerned with the second project: giving an account of responsible belief and epistemically responsible belief in particular rather than morally or prudentially responsible belief (although my account touches on that as well). I think this deserves attention on its own, no matter how it is related to the first project, that of epistemically justified belief. As I argue in the Appendix of the book, though, I think such epistemically responsible belief is a *necessary condition* for knowledge, so that the two projects turn out to be related after all.

Let me make one more preliminary comment before I address Goldberg's criticisms. According to Goldberg, my "main thesis is that we are responsible for a variety of actions that we perform, as well as the omissions, in the course of forming and sustaining belief" [p. 154, this volume]. That is true, but, crucially, that is not the whole story. I argue that we are *originally* responsible for a variety of actions and omissions in the course of forming and maintaining a belief (belief-forming actions and omissions), but that we are also *derivatively* responsible for our beliefs in virtue of that original responsibility.

Now, let us turn to Goldberg's criticisms. The first criticism concerns so-called doxastic compatibilism, that is, the idea (which I argue

against) that doxastic responsibility can be explained in terms of reasons-responsiveness (for this view, see, for instance, Hieronymi (2005), (2006), (2008); and McHugh (2014). According to Goldberg, doxastic compatibilism has at least two advantages over my Influence View.

First, on the reasons-responsiveness account, but *not* on my account, the subject is directly responsible for the belief itself. I am not sure why that is supposed to be an advantage of the reasons-responsiveness view. On my account, people are *derivatively* responsible for their beliefs and originally responsible for intellectual actions and omissions. They are still *responsible* for their beliefs: we can praise or blame them for their beliefs. We praise or blame people for all sorts of outcomes, like scientific discoveries, bomb explosions, a 1-year old child getting burned if we leave her close to a fire without keeping an eye on her, and so on. Derivative responsibility is as much responsibility as original responsibility; they are simply two different ways of being responsible [see p. 117 of the book].

Second, according to Goldberg, the reasons-responsiveness account can maintain an intimate connection between the verdict that a belief is responsible and the verdict that it is justified, whereas my Influence View on responsible belief cannot. Here is how Goldberg illustrates this claim. Imagine that Susie has a glitch in her cognitive system such that normal people, if they perform action A , come to believe that p on the basis of good evidence, whereas Susie, if she performs A , comes to believe another proposition q , and on the basis of bad evidence. Moreover, there is another action A^* that Susie could perform such that if she did that, she would come to believe that p on good evidence, but she blamelessly regards A^* as irresponsible. From her subjective point of view, what she should do is perform A . Hence, she performs A and comes to believe that q on bad evidence. My account implies that she acts responsibly and believes responsibly, for she is responsible for her beliefs and blameless. She has met her intellectual obligations or to the extent that she has violated them she is excused by blameless ignorance. Here's Goldberg's assessment of the case: "But surely – surely! – this is not a notion of responsible belief that has a significant role to play in a theory of epistemic justification or knowledge. At a minimum, it isn't a notion that is happily *identified* with epistemically justified belief" [p. 161, this volume].

I agree that my account of responsible belief does *not* provide a sufficient condition for knowledge. But, for all we know (and I defend this position in the Appendix of the book), it may still provide a *necessary* condition for knowledge: Susie would not have had knowledge if she had acted *contrary* to what she had believed was the epistemically responsible

thing to do. Plausibly, there are further conditions for knowledge, such as believing in accordance with one's evidence or at least not contrary to that evidence, that Susie does not meet. All I suggested is that if epistemic justification is a necessary condition for knowledge, it seems that epistemically responsible belief meets that condition. Moreover, the same applies to a reasons-responsiveness account. If Susie does *not* have the cognitive glitch, so that she is perfectly reasons-responsive, but culpably fails to meet an epistemic intellectual obligation to gather evidence that is relevantly related to p , she does not know that p either. This is because she fails to believe epistemically responsibly. Both the reason-responsiveness account and my influence account of responsible belief seem to provide a necessary condition for knowledge and they are therefore on a par when it comes to the issue of epistemic justification.

Now, one may object that there are no such things as *epistemic* intellectual obligations, since *epistemic reasons* cannot count in favor of actions, but only in favor of things that are true or false, such as beliefs. I realize that the notion of *epistemic* intellectual obligation is controversial. In the book, therefore, I address this objection in detail and argue that there *can* be epistemic reasons for actions and that there can be epistemic obligations to perform or not perform certain actions. Among other things, I point out that epistemic reasons can count in favor of suspending judgment. Suspending judgment, though, cannot be true or false, so epistemic reasons can count in favor of things that are neither true nor false [pp. 110-111].

Let us now turn to the second worry that Goldberg has. He suggests that my account of intellectual obligations implies that people have obligations that they can't possibly meet. Here is an example that he gives to illustrate this point. Susie can perform belief-influencing action A that is subjectively good and objectively bad. However, she can also perform another belief-influencing action A^* that is subjectively bad but objectively good: performing A^* would rid her of some beliefs that are based on insufficient evidence, but she has good reason to think otherwise. My account implies that she has a *subjective* obligation to do A and an *objective* obligation to do A^* . But she cannot do both. In conversation with Goldberg, I replied to this objection that Susie simply has a subjective obligation to do A , which she meets, and an objective obligation to do A^* , which she violates but for which she is excused by blameless ignorance.

Now, Goldberg's reply to my response to his objection is twofold. First, it "should strike us as peculiar indeed if one's account of intellectual obligations and excuses implies that there can be subjects of whom it is true both that (i) they are subject to intellectual obligations, and yet (ii)

it is in principle impossible for them to satisfy any of those obligations, as the obligations themselves systematically make conflicting demands of them” [p. 163, this volume]. In reply, let me point out that people don’t have cognitive glitches all the time. So, sometimes they are excused for not meeting an obligation. In the scenarios sketched, the subject can always meet at least one of the two conflicting obligations. If there’s an individual who has cognitive glitches everywhere, one may rightly question whether that person is subject to intellectual obligations at all. As I point out in *Responsible Belief*, people with cognitive glitches everywhere, such as people suffering from severe Alzheimer, are exempt from my analysis [p. 183]. Such people are not at all responsible for their beliefs. Moreover, there is nothing peculiar about the idea of being subject to conflicting obligations every now and then. Every meta-ethical theory that acknowledges the existence of both *subjective* and *objective* obligations (and that will be the vast majority of meta-ethical theories) implies that there can be situations in which what one *should do* deviates from what one *believes* one should do and, thus, that there can be conflicting obligations.

The second part of his reply is that my resulting account of intellectual obligations – one in terms of subjective and objective intellectual obligations – cannot be relevant for epistemic justification. I disagree: as I argue in the Appendix of the book, epistemically responsible belief, cashed out in terms of whether or not one has met one’s subjective and objective intellectual obligations, may well be *necessary* for knowledge, much like reasons-responsiveness. Again, I think the influence account of epistemically responsible belief stands, no matter how it relates to epistemically justified belief and knowledge, but for all we know, epistemically responsible belief (or something very closed to it) is necessary for knowledge. Thus, even though Goldberg’s examples show that epistemically responsible belief that is also true is insufficient for knowledge, my account may well shed important light on epistemic justification, for all we know as much as a reasons-responsiveness account.

*Department of Philosophy
Vrije Universiteit Amsterdam
De Boelelaan 1105
1081 HV Amsterdam, The Netherlands
Email: mail@rikpeels.nl*

ACKNOWLEDGMENTS

For their helpful comments on (the relevant parts of) this paper, I would like to thank Valentin Arts, Sandy Goldberg, Christopher Kulp, and Neil Levy. Sincere thanks are due to Luis Manuel Valdés-Villanueva for his kind and meticulous editorial work and helpfully guiding the whole process of bringing these articles into existence. This publication was made possible through the support of a grant from the Templeton World Charity Foundation. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the Templeton World Charity Foundation.

REFERENCES

- ALSTON, W. P. (1989), “The Deontological Conception of Epistemic Justification”; in *Epistemic Justification: Essays in the Theory of Knowledge*; Ithaca: Cornell University Press, pp. 115-152.
- BAYER, B. (2015), “The Elusiveness of Doxastic Compatibilism”; *American Philosophical Quarterly* 52.3, pp. 233-252.
- BOOTH, A. R. (2014), “On Some Recent Moves in Defence of Doxastic Compatibilism”; *Synthese* 191.8, pp. 1867-1880.
- COHEN, L. J. (1992), *An Essay on Belief and Acceptance*; Oxford: Clarendon Press.
- GOLDBERG, S. C. (2018), *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*; Oxford: Oxford University Press.
- HIERONYMI, P. (2005), “The Wrong Kind of Reason”, *The Journal of Philosophy* 102.9, pp. 437-457.
- (2006), “Controlling Attitudes”, *Pacific Philosophical Quarterly* 87.1, pp. 45-74.
- (2008), “Responsibility for Believing”; *Synthese* 161.3, pp. 357-373.
- JAMES, W. (1897), “Is Life Worth Living?”; in *The Will to Believe and Other Essays in Popular Philosophy*, Dover: Longmans, Green & Co. Johnston, 1956.
- LEVY, N. (2011), *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*; Oxford: Oxford University Press.
- (2014), *Consciousness and Moral Responsibility*; Oxford: Oxford University Press.
- (2016), “Dissolving the Puzzle of Resultant Moral Luck”; *Review of Philosophy and Psychology* 7.1, pp. 127-139.
- MCHUGH, C. (2014), “Exercising Doxastic Freedom”; *Philosophy and Phenomenological Research* 88.1. 1-37.
- MEYLAN, A. (2013), *Foundations of an Ethics of Belief*, Berlin: Walter de Gruyter.
- NOTTELMANN, N. (2007), *Blameworthy Belief: A Study in Epistemic Deontology*; Dordrecht: Springer.
- PEELS, R. (2011); “Tracing Culpable Ignorance”; *Logos and Episteme* 2.4, pp. 575-582.
- (2015). “Believing at Will Is Possible”, *Australasian Journal of Philosophy* 93.3, pp. 524-541.
- (2017), *Responsible Belief: A Theory in Ethics and Epistemology*; New York: Oxford University Press.

- RYAN, S. (2003), "Doxastic Compatibilism and the Ethics of Belief"; *Philosophical Studies* 114.1-2, pp. 47-79.
- STEUP, M. (2017), "Believing Intentionally"; *Synthese* 194.8, pp. 2673-2694.
- WAGNER, V. (2017), "On the Analogy of Free Will and Free Belief"; *Synthese* 194.8, pp. 2785-2810.