

# The Possibility of Love Independent Reasons

*Jussi Suikkanen*

Final Author Copy: to be published in *Essays in Philosophy*.

## 1. Frankfurt and Reasons

In both books *The Reasons of Love* and *Taking Ourselves Seriously & Getting It Right*, Harry Frankfurt argues that the practical reasons of an individual are grounded on what she loves (Frankfurt 2004 and 2006). This would entail that no agent could have practical reasons which would be independent of the objects which she happened to love.<sup>1</sup>

Frankfurt himself states these theses about reasons in the following quotations: “The things that are important to a person despite the fact that he does not actually care about them, or even know about them, can have importance to him only in virtue of standing in a certain relationship to something that he does care about (Frankfurt 2004: 22).”; “It is by caring about things that we infuse the world with importance (Frankfurt 2004: 23).”; “Love itself is a source of reasons (Frankfurt 2004: 37).”; “Love is the originating source of terminal value (Frankfurt 2004: 55).”; “Insofar as love is the creator of both inherent or terminal value and of importance, then, it is the ultimate ground of practical rationality (Frankfurt 2004: 56).”; “Authority of practical reason is based on the authority of love (Frankfurt 2006: 3).”; “When does a fact give us a reason for performing an action? It does so when it suggests that performing the action would help us reach one or another of our goals (Frankfurt 2006: 11).”; “In my view, it is only in virtue of what we actually care about that anything is important to us (Frankfurt 2006: 20).”; “There can be no rationally warranted

---

<sup>1</sup> Like Frankfurt, I use the term ‘objects’ broadly so as to include all potential objects of love: material objects, persons, processes, actions, works of art, abstracts objects like mathematical proofs and scientific theories, relationships, and so on (Frankfurt 2004: 41).

criteria for establishing anything as inherently important (Frankfurt 2006: 22).”; “Nothing is truly important to a person unless it makes a difference that he actually cares about. Importance is never inherent. It is always dependent upon the attitudes and dispositions of the individual (Frankfurt 2006: 23).”; “We possess an organised repertoire of final ends. That puts us in a position to determine...what we have reason to do (Frankfurt 2006: 28).”; “This is the doctrine of normative realism. It holds that there are objective reasons for us to act in various ways, whether we know them, or care about them, or not... My view is different. I do not believe that anything is inherently important... The standards of volitional rationality and of practical reason are grounded... only in ourselves (Frankfurt 2006: 33).”; “[L]ove is a powerful source of reasons... Insofar as a person loves something, he necessarily counts its interests as giving him reasons to serve those interests... Loving thus creates reasons by which the lover’s acts and devotion to this beloved are dictated and inspired (Frankfurt 2006: 42).”; “Through loving, then, we acquire final ends to which we cannot help but being bound; and by virtue of having those ends, we acquire reasons for acting that we cannot help but regard as particularly compelling (ibid.); and “Wholehearted love definitely settles, for each of us, issues concerning what we are to care about (Frankfurt 2006: 51).”

My interpretation of these passages (that they express a view according to which all reasons of an agent are grounded on what she loves) is shared by both Christine Korsgaard and Niko Kolodny. According to Korsgaard, “Frankfurt thinks it is only the things that we care about that give us reasons to act (Korsgaard 2006: 71)”, and, according to Kolodny, “[i]n the event, however, Frankfurt asserts something broader. Not simply are there reasons of love, he proposes, but all reasons are reasons of love (Kolodny 2006: 46).”

In order to properly understand Frankfurt’s view about reasons, we must consider his understanding of love. Frankfurt approaches love from the attitude of caring. For him,

caring about an object is ultimately a form of desiring (to have the object, for things to go well for the object, and so on) (Frankfurt 2004: 11). What is special about caring about something as a form of desiring is that, not only does one have the given desire, but one also wants that desire to be sustained (Frankfurt 2004: 16). This is the way in which, when one cares about something, one willingly grants importance to the desired object by committing oneself to desiring it.

According to Frankfurt, loving is then the most important mode of caring (Frankfurt 2004: 31).<sup>2</sup> Love, as an attitude, has several features which distinguish it from the lesser forms of caring (Frankfurt 2004: ch. 2). Firstly, loving something consists of continuously taking the interests of the loved object to constitute sufficient grounds for acting, and of being moved by these interests without a further thought (Frankfurt 2004: 37). Secondly, love must be disinterested. This means that the lover must promote the interests of the beloved for their own sake rather than instrumentally as a means to some further end (Frankfurt 2006: 40). Thirdly, love is the only form of caring such that its objects acquire value *because* we love them (Frankfurt 2004: 39). This is the case even if love is not typically a reaction to features of the objects which we judge to be valuable (ibid.).

Fourthly, love is always an attitude towards specific individual objects rather than being a generic concern for the good of whatever objects happen to fit a certain general description (Frankfurt 2004: 44). And finally, according to Frankfurt, what we love is not under our direct and immediate control (ibid.). The claim then is that what we care about in the important *loving* way, specified by the previous five conditions, determines what practical reasons we have (Frankfurt 2004: 59–62).

---

<sup>2</sup> This is why Frankfurt claims that love is a volitional attitude rather than anything affective or cognitive (Frankfurt 2004: 42).

Practical reasons are in this context understood to be considerations which count in favour of one's desires, intentions and/or actions (see e.g. Scanlon 1998: chap. 1). Thus, the view which Frankfurt has put forward in the quotations above can be summarised as the thesis that no consideration can count in favour of one's practical attitudes and/or actions unless that consideration is closely related to the promotion of the interests of an object which one cares about in an "involuntary, nonutilitarian, rigidly focused, and self-affirming" way (Frankfurt 2006: 40).

Of course, Frankfurt does not just state this view but rather he gives an interesting argument for it. That argument will be the main focus of this article. However, at this point, I need to ask patience from my readers. It will not be before the *next* section that I will describe in detail the argument that leads Frankfurt to believe that all practical reasons must be based on what we actually happen to love.<sup>3</sup> The rest of this article is dedicated to investigating and criticising that argument. My aim is to show that Frankfurt's ambitious argument fails to establish the conclusion that there are no 'love independent' reasons.

However, before I will describe Frankfurt's argument, I want to focus in the rest of this section on the previous conclusion that no agent can have reasons which would be independent of what she loved.<sup>4</sup> Let us first consider what kind of consequences this theory of reasons would have for the reasons of a potential rapist.

Frankfurt's view can easily explain what reasons a rapist would have for raping his victim.

During the act of raping, the rapist would get sexual gratification. This gratification would

---

<sup>3</sup> Frankfurt himself presents this argument in (Frankfurt 2004: 23–26) and (Frankfurt 2006: 22–25). I will also provide more detailed references and quotations of this argument in the next section in which I explain it. Very roughly, the basis crux of the argument is that any attempt to show that an agent had reasons which were independent of what she loved would inevitably be viciously circular (Frankfurt 2004: 24–25, and Frankfurt 2006: 23).

<sup>4</sup> I want to emphasise again that Frankfurt expresses this conclusion in the quotations which I gave in the beginning of this introduction, and that I share the interpretation of these conclusions with Korsgaard and Kolodny.

presumably be something that the rapist cares about involuntarily, rigidly, and for its own sake (see Frankfurt 2004: sec. 3 for Frankfurt's account of self-love and the egoistic reasons it grounds). As a result, given that caring about any object (even about sexual gratification) in this way is a source of reasons, the rapist would as a result have some reasons to rape his victim.

Would the rapist also have reasons not to rape his victim on Frankfurt's view? That view would ascribe such reasons to the rapist only if either one of two conditions were satisfied (Korsgaard 2006: 71). Firstly, the rapist might love his potential victim as the specific person she is. This would require that he cared about his victim's interests in the disinterested, nonutilitarian, involuntary, and rigid way. If the rapist loved his victim like this, then Frankfurt's view would ascribe him a reason not to rape her. Unfortunately, the potential victim could be a total stranger to the rapist, or even a person whom he hates. In these cases, the love of the particular person could not ground the rapist's reasons not to rape because the rapist would lack the required loving attitude.

Secondly, the rapist might love a general abstract moral ideal. Frankfurt understands morality to consist of a particular vision of "how [we are] to conduct ourselves in relations with other people", and "how our attitudes and actions should take into account the needs, the desires, and the entitlements of other people" (Frankfurt 2004: 7; and Frankfurt 2006: 28). Presumably, this vision of interpersonal relations would rule out the action of having sexual intercourse with a person against his or her will. If the rapist then cared about the fulfilment of the previous kind of a moral ideal involuntarily, rigidly, and for its own sake, then Frankfurt's view would again ascribe reasons to him not to rape his victim.

Unfortunately, as Frankfurt recognises, not everyone cares at all about the promotion of any abstract moral ideals of how persons are to be related to one another (Frankfurt 2004:

8–9, 71, and 98; Frankfurt 2006: 38, and 48; see also Bratman 2006: 81–84). Some agents might even love what is bad and evil, that is, human relations that would be the opposites of the ones characterised by the standard moral ideal. Thus, if our rapist did not love the moral ideal of how people are to be related to one another universally, the love of the abstract moral ideal could not be the source of his reasons not to rape either.

This means that one consequence of Frankfurt's view is that a rapist who does not love either his victim personally or an abstract moral ideal of certain kinds of interpersonal relationships does not have any reason not to rape his victim. In fact, given that he loves himself and his sexual gratification, it would be important for him to rape her.

I believe that a rapist has sufficient reasons not to rape his victim *no matter* what he happened to love. I am at least as certain about the truth of this belief as I am about the truth of the belief that I have hands.<sup>5</sup> The consequence of this is that Frankfurt's argument (explored in the next section) – the conclusion of which conflicts with my fundamental belief about reasons – would need to have premises that were more certainly true than my belief about the rapist's reasons. Otherwise, one would be warranted to use the argument's conclusion's awkward consequences to formulate a *reductio ad absurdum* argument against Frankfurt's theory about reasons.

This means that, if Frankfurt's argument premises are not all certainly true, then the conclusion of his argument and its consequences just provide reasons for us to believe that one his premises is false. Given that it is unlikely that all his premises are beyond doubt, there seems to be good reasons to believe that his argument fails to be sound even if it happened to be valid.

---

<sup>5</sup> In this sense, my basic belief captures a Moorean fact about reasons (see Lewis 1999 on Moorean facts).

At this point, it is worthwhile to pause both to consider a couple of responses which Frankfurt himself might give to the previous objection, and to locate Frankfurt's theory more broadly in the context of other Humean theories of practical reasons.<sup>6</sup>

(i) *Innately Loved Objects*. Firstly, Frankfurt could argue that it is part of the human nature that humans necessarily love certain objects innately. He could then try to argue that these loving attitudes will inevitably ground reasons for all human agents not to rape other human beings. In fact, Frankfurt endorses in many passages several innate loves (see Frankfurt 2004: 27, 47, and 71; and Frankfurt 2006: 35–38). I furthermore agree with him that there are at least some universally loved objects.

However, it is not clear that the fact that there are several things that human beings love innately suffices to guarantee that a rapist has reasons not to rape his victim. Frankfurt provides the following list of innately loved objects: one's own survival, avoiding crippling injury and illness, maintaining minimal contact with other human beings, being free from chronic suffering and endlessly stupefying boredom, being intact and healthy, being satisfied and in touch, our children, friends and other people close to one (ibid.). The problem is that, if we imagine a rapist who loves these objects innately, it does not necessarily follow that the interests of the objects he thereby loves are harmed in any way by his act of raping a stranger. He will still be in good health and satisfied, and so will be his family and friends.

The only way in which innately loved objects could help to avoid the objection would be if everyone innately loved every other human being, or an abstract, universal moral ideal about certain kind of human relations. This would guarantee that all potential human rapists would have reasons not to rape their victims even according to Frankfurt's view

---

<sup>6</sup> I thank the anonymous referee of *Essays in Philosophy* for pointing this out.

about reasons. Because we know that there are some actual individuals who do not care about the interests of everyone else, nor about any abstract moral ideals (Bratman 2006: 80), this response seems like a non-starter.

(ii) *Reasons and Morality*. Second, it could be pointed out that Frankfurt's conclusion only concerns an agent's reasons and what is important to him. This does not entail that moral standards would not apply to people like the rapist discussed above. Of course the moral standards apply to him wholly independently of what he happens to love (Frankfurt 2006: 46–47). There is no escape of the fact that the rapist's actions are wrong, bad, evil, wicked, morally forbidden, and so on no matter what his motivations are. Furthermore, given that he is motivated by his actual loves to rape his victim, we can also conclude from this that his character constituted by his loves is equally bad, wicked, and evil (Frankfurt 2004: 67). Thus, not only could the rapist be considered not to be a moral agent, but it would also be appropriate to think that he is positively an immoral agent.

Thus, on Frankfurt's view, whether moral standards apply to a rapist does not depend on what the rapist loves. Frankfurt also has an appealing explanation for this. The rapist is bound to harm a person or a moral ideal we love. This will make us angry. We will then express this anger by attributing moral blame to the rapist by using terms such as bad and evil (Frankfurt 2006: 47). Furthermore, we consider the moral distinctions to be objective (i.e., independent of what agents' love), because they are based on the need to protect objects which we could not conceive ourselves not loving (Frankfurt 2006: 46–47).

So, it might seem like Frankfurt has offered an appealing picture of morality which does not suffer from any of the problems of his view about reasons, importance, and inherent value. Of course, as such, this theory of morality does not make any of the previous Frankfurt's commitments about the rapist's reasons any more plausible. We can now



acknowledge that the rapist does something wrong and evil, and that he is a bad person. However, Frankfurt is still committed to saying that the rapist has *no reasons* to act in any morally more appropriate way. As far as the rapist is concerned, this just isn't important. So, the implausibility of Frankfurt's view about reasons remains even after we consider what he says about morality.

Furthermore, Frankfurt's view about reasons may also undermine his position on morality (Korsgaard 2006: 55–56). The problem is that Frankfurt explicitly states that “[e]ven if it were entirely clear what the moral commands, it would remain an open question how important it is for us to obey those commands (Frankfurt 2006: 28; see also Frankfurt 2004: 9).” In this sense, Frankfurt believes that morality itself is not normative, i.e., a source of reasons (Frankfurt 2004: 9). Depending on what an agent loves (as in the case of our rapist above), he or she might have no reasons at all to conform to the moral standards. This threatens to make all moral requirements very hollow indeed.<sup>7</sup>

*(iii) The Wider Context.* It is worthwhile also to set this objection into a wider context of the previous debates about Humean theories of reasons. Bernard Williams famously defended a theory of reasons which had very similar consequences as Frankfurt's view. On his view, an agent has a reason to do some particular act if and only if she either (i) already has a motive in her ‘subjective motivational set’ which would be served by doing that act, or (ii) there is a “sound deliberative route” from her prior motivations to such a motive (Williams 1981: 101–105; Williams 1995: 35–40). The sound deliberative route from the pre-existing motives to new ones refers here to what motives the agent could come to have by becoming more informed and more coherent in her desiring.

---

<sup>7</sup> Frankfurt himself explicitly accepts this conclusion but he does not believe that it is a problem: “[m]orality can provide at most only a severely limited and insufficient answer to the question of how a person should live (Frankfurt 2004: 7).”

Williams was explicit about the consequences of this view. He recognised that it will not offer everyone reasons to do, for instance, what is standardly considered to be morally required. In a famous case, he considered a husband who we believe should be nicer to his wife, but who has nothing in his motivational set that could move him to do so even after careful deliberation (Williams 1995: 39–40). Williams explicitly states that, according to his view, it would not make sense in this case to say that the husband has a reason to be nicer to his view. Yet, just like Frankfurt, Williams believes that this awkward conclusion is tenable because we can still call the husband “ungrateful, inconsiderate, hard, sexist, nasty, selfish, brutal, and many other disadvantageous things (Frankfurt 1996: 39).”

Many found Williams’s theory of reasons implausible just because it had the previous kind of consequences about what reasons bad people would have. Partly because of this, they began to look for flaws in Williams’s arguments for his view, and to explore alternative views about reasons and rationality (see, for instance, Scanlon 1998: 363–373, McDowell 1995, and Korsgaard 1986: 19–23).

More recently, Mark Schroeder has attempted to formulate a Humean theory of practical reasons which would generate enough reasons for everyone to avoid the previous kind of problems (Schroeder 2007). On his view, some fact is a reason for an agent to do a given act when the agent has some desire such that the fact’s obtaining would in part explain why the doing of that act would promote the satisfaction of that desire (Schroeder 2007: 59). Given that Schroeder sets no constraints on the content or nature of the relevant desires, this view is very liberal about what reasons agents have. Schroeder hopes that this element of his view would generate necessarily everyone reasons not to do immoral acts (Schroeder 2007: ch. 6).

So, for instance, imagine that the rapist would happen have a weak dispositional desire to know what has happened to his classmates from the primary school, and that one of them would appear in the evening news at 8 o'clock. Imagine also that, if the rapist committed the rape, he would miss his train home and so he would not watch the 8 o'clock news as he usually does. For Schroeder, this would be enough to generate a reason for the rapist not rape his victim. After all, there would be some fact about the rape that would in part explain why his act of not committing the rape would promote the satisfaction of one of his desires (namely, the one to know about his classmates). The hope is that, because all agents have huge sets of desires for almost infinitely many objects, we could always generate similar reasons not to rape others for all potential rapists.

It is instructive to note why this strategy of dealing with the rapist's reasons is not available for Frankfurt. He explicitly argues against the idea that the promotion of any desire whatsoever could ground reasons (Frankfurt 2006: 10–11). This is because many of our often strange desires may seem from our own perspective like alien forces happening to us. Because of this, we externalise these desires by not permitting them to guide our behaviour. By doing so we deny that their satisfaction would make a significant difference to our lives, and so we reject the idea that these desires could ground our reasons. This is why Frankfurt believes that only the attitude of love can generate reasons (Frankfurt 2006: 11–14, and 24–26). Given that we love far fewer things than we desire, it is much harder for Frankfurt to argue that everyone's loving concerns generate reasons not to rape others than it is for Schroeder to argue that everyone must have at least some desire which will be promoted by not raping.

Thus, to summarise, Frankfurt's position has the same awkward consequences about our reasons for not doing immoral actions as Williams's view. These positions are thus implausible for the very same reasons. It is true that Mark Schroeder has offered a

potential Humean way to avoid these problems, but it is clear that his view is not available for Frankfurt. It would require giving up the idea that only *love* can ground reasons.

Besides the rape case, the conclusion of Frankfurt's argument has also another strange consequence. Frankfurt explicitly states that, if one did not love anything, nothing would be important for one, and thus one would not have any reasons to do anything (Frankfurt 2004: 22, 26, and 58; Frankfurt 2006: 20 and 23). Reasons and importance thus come to the scene only after we begin to love objects. Hence, Frankfurt states that by loving objects 'we infuse the world with importance' (Frankfurt 2004: 23). If there were antecedent reasons for us to love objects before we loved them, then objects in the world would have importance with respect to us independently of whether we happen to love them. Because Frankfurt denies this, he must believe that, when we happen to begin to love an object, we create reasons for ourselves (Frankfurt 2006: 25). When we begin to love objects, considerations which before were not reasons for us would begin to count in favour of our actions.

Many philosophers believe that it is implausible that we would have this kind of a power to create reasons by adopting attitudes towards objects.<sup>8</sup> Furthermore, we can also ask, what reasons could we come to have as a result of loving objects which we did not have any antecedent reasons to love (Korsgaard 1997, Dancy 2000: 26–43)? Intuitively, we would need 'reasons in' (to love a given object) in order to get 'reasons out' (to act on the basis of

---

<sup>8</sup> This idea that we could create reasons by beginning to love things is sometimes called 'bootstrapping' (Bratman 1987: 24–27). It is usually considered to be a problematic feature of the views according to which creating reasons by adopting attitudes is possible. Frankfurt thinks that a bootstrapping objection can be made against the view which begins from the idea that there are reasons to love objects, and by finding out what these reasons are we can rationally begin to love new objects (Frankfurt 2004: 26). His argument described below is supposed to show this. Because I think that that argument is false, it is hard to see how the alternative view could lead to objectionable bootstrapping.

Frankfurt might reply that the bootstrapping objection does not apply his view because of the fact that, for him, what we love is not under our voluntary control. This is why we cannot create reasons by deciding to love objects. However, this seems to make our power to create reasons by non-voluntarily beginning to love objects even more mysterious. Just why only an attitude which we cannot control could have that kind of a consequence?

that love). If Frankfurt will be unable to give a plausible answer to how we could get reasons on the basis of an attitude we have no reasons to have, then his view might fail to leave room for any practical reasons at all. So, either he has to give a further account of our mysterious ability to create reasons, or he will have problems explaining how we could have any reasons in the first place.

So far, I have merely drawn out certain implausible consequences of Frankfurt's argument's conclusion. The problem with arguing against him in this way is that, at best, it could only show that his argument must be either unsound, invalid, or both. However, using this argumentative strategy could never show just where, in which premise or a step, the fault lies in his argument. Learning this would be far more instructive.

To find out where his argument fails, I plan to explore an analogy between practical reasons for loving objects and theoretical reasons for beliefs.<sup>9</sup> In the next section, I begin by explaining Frankfurt's argument against the love independent reasons. I then sketch, in section 3, an analogical argument which could be made about our reasons for beliefs. As it happens, it will be easier to see where this analogical argument goes wrong. This will help us also to see more easily where Frankfurt's original argument fails (I will attempt to explain this in section 4). And, without this argument, Frankfurt gives us no reason to doubt our intuition that there are at least some practical reasons which are independent of what we love.

Finally, there are two things worth noting before we proceed. Firstly, even though the analogy between practical reasons and the reasons for beliefs turns out to be useful during this investigation of Frankfurt's argument, it might in the end be that we ought to give different accounts of these two kinds of reasons. However, even if this were the case, this

---

<sup>9</sup> Outside the context of Frankfurt's argument, this analogy has been extensively explored by Terence Cuneo (2007).

would not affect the problems of Frankfurt's argument explored here. Secondly, I will not attempt to prove that we do have practical reasons which do not depend on what we love. My only aim is to show that Frankfurt's argument does not establish that there are no such reasons. Perhaps there are other better arguments to the same conclusion.

## **2. Frankfurt's Vicious Circle Argument**

Frankfurt focuses on a situation in which I begin to wonder whether I ought to love an object which I have not loved before (Frankfurt 2004: 23). In this situation, I thus inquire whether I would have good reasons to begin to love a new object.<sup>10</sup> Frankfurt's argument is based on the claim that posing this question and trying to answer it immediately leads to a vicious circle (Frankfurt 2004: 24, and 26). According to him, '[n]o attempt to deal with the problem of what we have good reason to [love]—to deal with it systematically and from the ground up—can possibly succeed (Frankfurt 2004: 24).' This is a radical claim considering its consequences explained in previous section. But, just why would it be impossible to inquire what one ought to love by considering the reasons there are? Here is Frankfurt's explanation of why this would be the case (Frankfurt 2004: §10–§11, and 2006: 22–23).

Frankfurt's argument relies on an intuitively plausible assumption. According to it, the objects which one loves must necessarily have a significant effect on what one's life is like (Frankfurt 2004: 25). This is why activities such as counting blades of grass could never be objects of our love. This assumption entails that only the considerations which make a significant difference to what one's life is like could ever be reasons for beginning to love

---

<sup>10</sup> I assume that the reasons to love new objects are important because they could be claimed to give rise for all other practical reasons. In one way or another, all these other reasons would be instrumental reasons to try to obtain the objects our love, reasons to admire and enjoy them, reasons to try to protect and respect them, and so on (see Frankfurt 2004: 55). Thus, maybe there is a way to show that all ordinary reasons for actions and conative attitudes are in a way instrumental reasons derived from the basic reasons to love objects. However, if the reasons to love objects were independent of what we have loved before, then these other derivative reasons too would be independent of what we loved at a given moment. Because we would have antecedent reasons to love an object, we would also already have all the other reasons too independently of what we loved.

a new object and to act on such a love (ibid.). Hence, ‘inconsequential’ changes to what one’s life is like could not be sufficient to function as reasons to adopt new loves. Acting on the reasons for loving new objects thus needs to make an ‘important’ difference to what one’s life is like.

We can then ask, how could one identify which considerations would make an important enough difference to the kind of life one lives? This question gets us to the heart of Frankfurt’s argument. Frankfurt claims that, in trying to identify which differences between different lives which one could live would be significant enough, one needs to have already committed oneself to some standards for evaluating importance (Frankfurt 2004: 24; Frankfurt 2006: 23). And, he believes that such commitments to evaluative standards can only be grounded on what one already loves (ibid.). This means that inquiring what new objects one should love requires that one already loves other, related objects (Frankfurt 2004: 25–26; Frankfurt 2006: 23). One has reasons to adopt loving attitudes towards new objects only if, on the basis of these previous loves, one comes to judge that the new loved objects would make a significant enough difference to one’s life.

So, we can have reasons to adopt new loves only on the basis of the objects we already loved before the adoption of the new loves. Of course, this fact can often escape our attention from our first-personal perspective from which we reflect about what we should love. This is why any attempt to argue for reasons to adopt new loves that are not based on one’s previous loves are only implicitly viciously circular. This circularity can, however, always be exposed in further reflection. Furthermore, if we genuinely did not love anything, then we would be stuck. We would not have reasons to begin to love anything (Frankfurt 2004: 22, 26; Frankfurt 2006: 23–24). This is because we could not identify standards by which to assess whether the potential new objects of our love would make a significant enough difference to the way we lived.

Reasons therefore appear only after we have begun to love objects. What other objects we then find out we have reasons to love is a function of what we already loved. Thus, ultimately ‘the normative question of how one should live ... can sensibly be asked only on the basis of a prior answer to the factual question of what [one] actually [loves] (Frankfurt 2004: 26).’ We already saw in the previous section that this conclusion has many implausible consequences. What we then need to find out is just where the argument for it goes wrong.

At this point, it is worthwhile to say something about one reason why we should be suspicious about Frankfurt’s argument and its validity. We should notice that, at the crucial moment, it draws a normative (or a metaphysical) conclusion about the existence of reasons on the basis of epistemological considerations. The argument first observes that it is impossible to recognise whether some life-change is significant enough to ground a reason without relying on what one already loves.<sup>11</sup> It then concludes from this that the reasons which one has for adopting new objects of love must be a function of what one already happens to love.<sup>12</sup> So, the argument moves from psychological and epistemic constraints on our ability to recognise significance to what normative constraints there are on what reasons we have.<sup>13</sup>

---

<sup>11</sup> As Frankfurt puts it: “In order to know how to determine what is important to himself, then, he must already know how to identify certain things as making differences that are important to him (Frankfurt 2004: 25–26)”, and “The truth is, I believe, that it is possible to ground judgments of importance only in judgments concerning what people care about (Frankfurt 2006: 23)”.

<sup>12</sup> Frankfurt formulates the conclusion after the previous quotation in footnote 11 in the following way: “Nothing is truly important to a person unless it makes a difference that he actually cares about. Importance is never inherent (Frankfurt 2006: 23).”

<sup>13</sup> Michael Bratman voices similar concerns. According to Bratman, Frankfurt correctly claims that loving entails treating certain considerations as reasons and it can also prevent treating other considerations as reasons (Bratman 2006: 81–82). However, Bratman points out that there is a missing step from this psychological thesis to the normative claims about what reasons agents have (*ibid.*).



This step relies on the thought that there could not be facts about the significance of life-changes or about reasons which we would be unable to recognise.<sup>14</sup> We should be suspicious of whether this assumption is true. It has been recently forcefully argued that, in any domain in which we can talk about knowledge, there will be some truths which will be unknowable (Williamson 2000: ch. 4). If this is right and we can sometimes know which life-changes are significant and what reasons we have, then some facts about significance and reasons will be unknowable in any case. This would undermine the central step in Frankfurt's argument which assumes that facts about the significance of life-changes must be such that we are in a position to recognise them.

### **3. An Analogical Argument in the Context of Theoretical Reasons**

At this point, it is useful to consider how Frankfurt's argument would work in the case of theoretical reasons for belief. This new argument will at some stage begin to sound confused. This will give us some indication of where the original argument about practical reasons goes wrong.

Let us consider a situation in which I find myself with a limited set of beliefs about the world. I have formed these beliefs on the basis of correctly reasoning from my previous experiences. I have done so much reasoning that I could no longer form any new justified beliefs from them. I have exhausted both my existing empirical data and my potential for *a priori* reasoning. I then become curious. I want to find out whether there are any other beliefs which I ought to adopt. How should I answer this question?

One plausible suggestion is that I should begin by getting new experiences. These experiences could then serve as good reasons for adopting new beliefs. However, it could

---

<sup>14</sup> This thought is encapsulated in the sentence already quoted above: "Nothing is truly important to a person unless it makes a difference that he actually cares about (Frankfurt 2006: 23)." The constraint of caring is given by Frankfurt because caring enables the agent to treat some life-change as significant (ibid.).

be claimed that any ‘inconsequential experiences’ (those that, for instance, were mere imaginations, appearances, dreams, illusions or hallucinations, or even just too similar to my previous experiences) would not suffice. Therefore, the new experiences would need to be ‘significant enough’ in order to be able to function as reasons for adopting new beliefs.

The question then is, how could I identify which of my new experiences would be significant enough for being able to justify the new beliefs? In order to answer this question, let us try to follow Frankfurt’s reasoning in the context of practical reasons.

In this case, we should claim that, in identifying which new experiences would be significant enough to be reasons for new beliefs, I would already need to be committed to some standards for evaluating epistemic significance. It might then be natural to think that my old set of beliefs would constitute the standards which I would need for being able to assess the significance of the new experiences.

However, it could be pointed out at this point that my quest for adopting new beliefs is bound to end in a vicious circle. If I can assess whether some experience is significant enough to be a reason for a new belief, then I must already have beliefs about (i) how likely it is that the given belief would be true (on the basis of the evidence which I already had for it), and (ii) how likely it is that my experience is appropriately related to the object of the given belief so as to count as justification for it. If the new experience fails to satisfy these standards set by my old beliefs, I will only judge that the new experience is not significant enough to count in favour of the adoption of a new belief.

This would mean that, even if try to have new beliefs and think that these beliefs will be based on the new experiences as reasons, it still turns out that the new beliefs could only be insignificant extensions of my older beliefs. In principle, I could have already adopted

these beliefs on the basis of my previous beliefs, or at least the old beliefs significantly constrain which new beliefs I can come to adopt. And, without my earlier beliefs, I could not even begin the process of adopting new beliefs, because I could not assess the significance of my new experiences as reasons for beliefs. So, again, either I am faced with a vicious circle, or I cannot even begin to adopt new beliefs.

Hardly anyone believes that this really is our epistemic situation with respect to adopting new beliefs. Intuitively, we can adopt new justified beliefs which are not a direct consequence or a part of what we already believed. And, it seems appealing to think at least we can have reasons to adopt new beliefs even if our old beliefs can occasionally prevent us from recognising these reasons and thus block us from adopting the new beliefs.

Consider the classic example of the Japanese holdouts after the end of World War II. Some of these soldiers lived on the small Pacific islands until the early 1970s believing that Japan was still engaged in a war against the United States. Their sets of beliefs were so unified and entrenched that, no matter what evidence the Americans tried to provide for the soldiers about the end of the war, the holdouts still continued to believe that the war had not finished. The Japanese soldiers just interpreted all the new material as unreliable and easily dismissible enemy propaganda.

Therefore, given their earlier beliefs, the Japanese soldiers were bound to fail to recognise the epistemic significance of their new experiences. Yet, few of us would be willing to think that, because of this, the soldiers had *no reason at all* to begin to believe that the war had ended. We are inclined to think that the evidence offered to the soldiers was a reason for them to change their beliefs, even if their previous beliefs prevented them from seeing this normative fact. However, if the Frankfurtian argument explained above works, then

we should give up these intuitions. Fortunately, in this context, it is easier to see where the reasoning of that argument goes wrong.

#### **4. The Fault in the Argument**

We can then ask, where did the previous argument about the reasons for beliefs go wrong? The crucial flaw seems to be the assumption according to which the criteria for assessing the significance of the new experiences as reasons for new beliefs must be limited to the believer's own antecedent set of beliefs. As the case of the Japanese soldiers illustrates, there just seems to be no reason why we should accept this assumption given how bizarre the antecedent beliefs can be. The new experiences' ability to function as reasons for beliefs depends instead on whether they satisfy all the epistemic norms which are relevant for how justified beliefs are.

It is true that one epistemic norm which should be used in evaluating the epistemic significance of an experience is whether it and the potential beliefs which it might support cohere with the believer's prior beliefs and experiences. Of course, as a matter of fact, this often is the standard which the believer herself uses to assess her new experiences. As individual believers, we can hardly get outside our own realm of beliefs in deliberation.

However, that internal standard is not the only criterion which should be used to assess the significance of a believer's new experiences as reasons for beliefs. When the new experiences of a given believer are assessed, we can, for instance ask the following questions:

- Do other observers have similar experiences in the same situations?
- Do the new experiences and the beliefs based on them conflict with the justified beliefs and knowledge of other believers?
- Are the circumstances of the new experience prone to create distorted experiences?

- Have the new experiences been gained by using fully functioning perceptual faculties which rely on reliable causal mechanisms?
- Was the experiencer momentarily tired, drugged, inattentive, or the like?

All these considerations are external to the believer's own antecedent set of beliefs, and yet they are still relevant for assessing whether her new experiences are good reasons for her to adopt new beliefs.<sup>15</sup>

This means that the epistemic norms which determine whether a given experience is a good reason for a new belief need not be limited to the believer's own antecedent beliefs. Rather, they can also include all the other considerations that are relevant for whether the beliefs formed on the basis of a given experience are likely to be true. This entails that the believer herself may be sometimes unable to correctly identify whether her experiences really are significant enough justify the adoption of any new beliefs.

However, we should not think that, just because an agent fails to identify a new experience as a reason for a new belief, that experience could not be a reason for her. If an agent fails to recognise the significance of her new experiences, she just fails to form all the beliefs she has reasons to have. It seems to be a good feature of the sketched pluralist view about justification that it leaves room for this kind of epistemic irrationality. This is vividly illustrated by the case of the Japanese soldiers. At least intuitively, they were given many epistemic reasons which they failed to recognise.

With these epistemological insights in mind, we can return to Frankfurt's original argument about our practical reasons. We are now in a position to appreciate that his argument assumes that the criteria with which we should assess the significance of the

---

<sup>15</sup> This is the standard externalist picture of justification in epistemology. For both important defences of this view and good objections to the internalist alternatives, see Armstrong (1973), Dretske (1981), Nozick (1981: ch. 3), and Goldman (1986).

changes in an agent life must be limited to that agent's previous set of attitudes of love (Frankfurt 2004: 22).

If we look at the lesson which we learned from the previous discussion of the theoretical reasons, this is probably something which we should not assume. Instead, we should use our general normative standards to assess whether a given change to an agent's life would be significant enough to ground a reason for her to love new objects. It is true that *one* important criterion for the life-change is how the new life would fit what the agent has loved before. If the impact of the new potential object of love matches the agent's antecedent loves, then the way in which that object would change the agent's life would at least appear to be significant from the agent's own first-personal perspective.

Of course, this is the criterion which actual agents as a matter of fact tend to use for assessing what they should love. In the same way as we cannot get outside our beliefs as believers, we can hardly get outside our loves as agents and evaluators.

But, it is by no means obvious that the only question to ask in assessing the changes to an agent's life is whether the some new way of life would fit what she already loved. Firstly, from our third-personal perspective, we can ask the following kind of questions about the agent's antecedent loving attitudes:

- Are the agent's loving attitudes stable under her self-reflection?<sup>16</sup>
- Would the agent still deem the change in her life to be significant if her antecedent set of loving attitudes were made more informed, coherent, and unified?<sup>17</sup>

---

<sup>16</sup> This norm is emphasised by, for instance, Sidgwick (1884: 339–342), Scanlon (1998: 65–70), and Crisp (2006: 88–91).

<sup>17</sup> This norm is emphasised by, for instance, Smith (1994: sec. 5.9) and Williams (1981). Even Frankfurt acknowledges towards the end of his 2006 book that we can come to correct our views on what we are to love and thus what is important to us by attempting to make our loves more coherent and by becoming more informed about the objects of our love (Frankfurt 2006: 49). This is problematic in the light of his previous argument. If we consider the objects loved by hypothetical versions of ourselves whose loves have been

- Could the agent see the change in her life as significant also from the perspectives of other agents?<sup>18</sup>

If the agent's assessment of the significance of a life-change would be different after we had idealised her loves in these ways, then, in assessing the changes to her life, we should not give much weight to how significant something appears to be from the perspective of the agent's own actual loves. Rather, we should consider what the agent would make of these changes from the perspective of her improved loves. After all, this is how the agent would want to assess her life-changes herself too if she weren't suffering from irrationalities.

We can also ask other, more external evaluative questions about the agent's antecedent objects of love. We can, for instance, ask:

- Does the agent love objects which are worth loving?
- Does she love similar objects which other reasonable and informed people love?
- Does she love the good, the beautiful, the just, and the kind?
- Are the objects of her love such that loving those objects leads to happiness, well-being, and flourishing human life?
- Do her loving attitudes conform to the important ethical standards?<sup>19</sup>

---

corrected in these ways, it might be that we, as actual agents with the uncorrected loves, could not recognize the importance of those objects to us. If Frankfurt wants to hold onto the idea that actual agents must be in a position to recognize what is important to them, then this latter view of which acknowledges the role of the corrected loves as a source of reasons would be ruled out by that constraint.

Also, Frankfurt still believes that, even after the loving attitudes are corrected in these two ways, people will still love different objects and thus have different reasons (Frankfurt 2006: 50). Smith is more optimistic that already this process of idealizing the loving attitudes can get rid of that kind of relativism about reasons espoused by Frankfurt. Because the process can create new loves and get rid of old ones, it might be that in the end all rational agents converge on what they love in identical situations (Smith 1984: 173). This would mean that what they are to love and what is important to agents would not depend on the actual, uncorrected loving attitudes these agents happen have at any given moment.

<sup>18</sup> See Broome (2004: 42).

<sup>19</sup> Of course, these external standards for assessing what one should love are advocated by the standard Kantian, Aristotelian, and Intuitionist theories of practical reason. For a helpful overview of these views, see Cullity & Gaut (1997).

After this, we can consider whether the agent would judge some change in her life to be sufficiently significant to ground reasons for loving new objects, if her antecedent loves would conform to these external standards. If this were the case, we could claim that a change in an agent's life can be significant enough to ground reasons for her to love new objects even if she cannot actually recognise that change as such. This has the plausible consequence that an agent's judgments about the significance of life-changes and reasons can be fallible.<sup>20</sup>

Now, it could be objected that this only shows that, in assessing whether some changes to another agent's life would be significant enough for grounding reasons for her, we - the evaluators of her life and reasons - must rely on what we happen love at the moment of assessment.<sup>21</sup> Perhaps this is right, but it is questionable why this would be problematic. It would still be the case that, contrary to what Frankfurt claims, whether some changes in an agent's life are significant enough is not solely determined by what that agent loves herself. Furthermore, we can hardly be required to give up our own convictions about what is significant when we assess the life-changes of another person, and what reasons they have. Requiring that we would do so would be analogical to requiring that we should give up our beliefs when we assess the truth and the epistemic status of other people's beliefs.

However, our view should also leave room for the fact that our own evaluations about the significance of the changes to someone else's life can be mistaken. Fortunately, such room for our own mistakes exists in our view, because it can again be asked, for instance, what

---

<sup>20</sup> Frankfurt explicitly accepts that an agent can be mistaken about how significant some change to her life would be (Frankfurt 2004: 21). However, according to him, such mistakes are based of us failing to realize what we already love (Frankfurt 2004: 22).

<sup>21</sup> This certainly is the case for expressivists who think that evaluative language in general expresses our non-cognitive pro and con attitudes. Yet, expressivists strongly argue that their view does not have the kind of relativist consequences such that which changes in someone else's life are significant depend on what that agent happens to love (see, for instance, Blackburn (2006)). Instead, expressivists like Blackburn believe that we should stand by our own attitudes when making evaluations of other people's lives and their reasons (ibid.). Frankfurt seems to be unable to adopt this position because he believes that those other people would have to be in a position to recognize the reasons which we would want to ascribe to them on the basis of what they love (Frankfurt 2006: 23).



we would love ourselves if our loving attitudes were made more informed, coherent, and unified. It can also be asked what we would love if our current loves would be made to conform to the external standards with respect to what is worth loving.

If our own loves were idealised in these ways, it might be that we would judge the significance of the changes to other agents' lives and their reasons differently.<sup>22</sup> This shows that our actual judgments about these issues can be mistaken, also by our own lights.

It might, of course, be the case that we can only investigate what we would love in the idealised circumstances by starting our investigation from what we actually happen to love. However, even this would not show that the other agents' reasons for loving attitudes are actually constrained by what we happened to love. It only shows that *our beliefs* about their reasons might be constrained in that way. It is still an open question whether these beliefs are true or not. Our epistemic situation with respect to the practical reasons would then be akin to the classic Neurath's boat type of a situation. We can only come to make more accurate judgments about our own and others' lives and reasons by gradually correcting our loves one by one on the basis of our other loving attitudes.

We are then finally able to grasp where Frankfurt's influential argument goes wrong. He assumes that *the only standard* for assessing the significance of the changes to an agent's life is what the agent herself would judge to be significant on the basis of what she happens to love (Frankfurt 2004: 24–25; Frankfurt 2006: 23). However, there are no good reasons to believe that this is the only relevant standard.<sup>23</sup> This is because it is plausible that the

---

<sup>22</sup> Expressivists often also emphasize this fact (see Blackburn (1984: 197–202)). This is a natural way for the expressivists to capture our talk about the possibility that we are mistaken and our talk about ethical facts and truths within their theory without any additional metaphysical package.

<sup>23</sup> At least Frankfurt says nothing why other norms could not be relevant. Of course, there could be some good reasons for this. For instance, it could be argued that practical reasons should be capable of motivating

agent could be mistaken about which changes in her life would be significant in ways that go beyond failing to realise what she already loves.

It can thus also matter (i) which life-changes the agent would judge to be significant if her loving attitudes were corrected in many internal and external ways, (ii) whether the relevant changes in her life are significant in the light of the evaluative standards which we use, (iii) whether we would judge the relevant life-changes to be significant if our own loving attitudes were corrected in both internal and external ways, and (iv) whether the changes in the agent's life are significant by whatever external, objective standards of significance there might be.<sup>24</sup> If this is right, then some changes to an agent's life can be significant enough to ground reasons for her for beginning to love new objects independently of what the agent already loves.

## 5. Conclusion

I began from Frankfurt's argument. It tried to establish that which new objects we have reasons to love (and thus what practical reasons we have more generally on the basis of our loves) is determined by what we already happen to love. In the first section, I flagged out that this conclusion has many implausible consequences. It fails to leave room for certain intuitive moral reasons, and it allows 'bootstrapping' reasons into existence. Because of this, it has been worthwhile to investigate closely just where Frankfurt's argument makes a mistake.

---

the agent to action (Williams 1981). It could be then claimed that considerations that are based on life-changes which the agent could not see as significant given what she loves could not satisfy this condition for being a reason. Even if this were the case, this would only mean that Frankfurt's argument would collapse into Williams' much discussed and controversial standard argument for reasons-internalism.

<sup>24</sup> Note that my argument does not require that there are any external objective standards of significance (clause (iv)), or that any external corrections could be made to the agent's or to our own loving attitudes (elements in clauses (i) and (iii)). My objection to Frankfurt's argument succeeds without these additional elements – it would still be the case that which life-changes are significant is not a direct function of what an agent happens to love. For this reason, my objection to Frankfurt's argument neither begs the question against him by assuming objective standards of evaluation and universal reasons, nor requires any extravagant commitments. Conversely, if Frankfurt's argument is based on the assumption that there are no external standards to assess significance, then it threatens to beg the question against his opponent who believes in love independent reasons.

I did this by first sketching a corresponding argument, which, if successful, would show that which of our new experiences are theoretical reasons for adopting new beliefs is likewise determined by what we already happened to believe. However, this thesis about our epistemic reasons is more clearly false as illustrated by the case of the Japanese holdouts.

The epistemic norms which determine what reasons for adopting new beliefs there are need not be solely based on what beliefs we already have before we come to have new experiences. Instead, such norms can also be sensitive to a wider set of considerations which are relevant for whether the beliefs formed on the basis of different kinds of experiences are likely to be true. This already gives us reasons to doubt whether the corresponding standards for assessing what practical reasons a given agent has need to be constituted by what she already loves.

According to Frankfurt, the practical reasons of an agent (including the reasons to love objects) are a function of whether acting on those reasons would lead to significant enough changes in the agent's life (see, again, the quotations in the beginning). One relevant factor in assessing whether a given change is significant enough in an agent's life is to consider whether the life-change would appear to be significant from the perspective of the agent's own actual loves. However, contrary to what Frankfurt assumes, this is not the only relevant criterion by which to assess the significance of life-changes generally.

It is also relevant what the agent would judge to be significant if her antecedent loves would be corrected in the obvious ways, what we judge to be significant as evaluators of her life and reasons, what kind of judgments we would make if our own loves would be corrected, and whether the changes in the agent's life would be significant by whatever

objective criteria of significance there happens to exist. If we recognise these standards for assessing the significance of life-changes, then, contra Frankfurt, it is not true that what reasons an agent has is a direct function of what she already loves. Some other argument would have to be given to that conclusion.

This means that we can continue to believe that there are reasons that are independent of what the person who has these reasons loves. As a result, we can also rationally continue to believe that there are moral obligations, that the reasons which these obligations provide can be independent of what we love, and that we cannot bootstrap reasons into existence simply by starting to love objects.

### References

- Armstrong, David (1973): *Belief, Truth, and Knowledge*. Cambridge University Press, Cambridge.
- Blackburn, Simon (1984): *Spreading the Word*. Oxford University Press, Oxford.
- Blackburn, Simon (2006): “Must We Weep for Sentimentalism?”. In J. Dreier (ed.): *Contemporary Debates in Moral Theory*. Blackwell, Oxford, 144–160.
- Bratman, Michael (1987): *Intention, Plans and Practical Reason*. Harvard University Press, Cambridge, Mass.
- Bratman, Michael (2006): “A Thoughtful and Reasonable Stability – A Comment on Harry Frankfurt’s 2004 Tanner Lectures”. In H. Frankfurt: *Taking Ourselves Seriously & Getting It Right*. Stanford University Press, Stanford, 77–90.
- Broome, John (2004): *Weighing Lives*. Oxford University Press, Oxford.
- Crisp, Roger (2006): *Reasons & the Good*. Oxford University Press, Oxford.
- Cullity, Garrett & Gaut, Berys (1997): “Introduction”. In Cullity & Gaut (eds.): *Ethics and Practical Reason*. Oxford University Press, Oxford.
- Cuneo, Terence (2007): *The Normative Web*. Oxford University Press, Oxford.

- Dancy, Jonathan (2000): *Practical Reality*. Oxford University Press, Oxford.
- Dretske, Fred (1981): *Knowledge and the Flow of Information*. MIT Press, Cambridge, Mass.
- Foot, Philippa (1972): "Morality as a System of Hypothetical Imperatives". *Philosophical Review* 81, 305–316.
- Frankfurt, Harry (2004): *The Reasons of Love*. Princeton University Press, Princeton.
- Frankfurt, Harry (2006): *Taking Ourselves Seriously & Getting It Right*. Stanford University Press, Stanford.
- Goldman, Alvin (1986): *Epistemology and Cognition*. Harvard University Press, Cambridge, Mass.
- Harman, Gilbert (1975): "Moral Relativism Defended". *Philosophical Review* 84, 3–22.
- Kolodny, Niko (2006): "Review of *Reasons of Love*". *Journal of Philosophy* 103, 43–50.
- Korsgaard, Christine (1986): "Scepticism about Practical Reason". *Journal of Philosophy* 83, 5–25.
- Korsgaard, Christine (1997): "The Normativity of Instrumental Reason". In Cullity & Gaut (eds.): *Ethics and Practical Reason*. Clarendon Press, Oxford, 215–254.
- Korsgaard, Christine (2006): "Morality and the Logic of Caring – A Comment on Harry Frankfurt". In H. Frankfurt: *Taking Ourselves Seriously & Getting It Right*. Stanford University Press, Stanford, 55–76.
- Lewis, David (1999): "Elusive Knowledge". In his *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 418–445.
- McDowell, John (1995): "Might There Be External Reasons?". In Altham & Harrison (eds.): *World, Mind, and Ethics – Essays on the Moral Philosophy of Bernard Williams*. Cambridge University Press, Cambridge, 68–85.
- Nozick, Robert (1981): *Philosophical Explanations*. Oxford University Press, Oxford.
- Scanlon, Thomas (1998): *What We Owe to Each Other*. Belknap Press, Cambridge, Mass.
- Schroeder, Mark (2007): *Slaves of the Passions*. Oxford University Press, Oxford.

- Sidgwick, Henry (1884): *The Methods of Ethics*, 3rd ed.. MacMillan and Co., London.
- Smith, Michael (1994): *The Moral Problem*. Blackwell, Oxford.
- Smith, Michael (1995): “Internal Reasons”. *Philosophy and Phenomenological Research* 55, 109–131.
- Williams, Bernard (1981): “Internal and External Reasons”. In his *Moral Luck*. Cambridge University Press, Cambridge, 101–113.
- Williams, Bernard (1995): “Internal Reasons and the Obscurity of Blame”. In his *Making Sense of Humanity – and Other Philosophical Papers*. Cambridge: Cambridge University Press, 35–45.
- Williamson, Timothy (2000): *Knowledge and Its Limits*. Oxford: Oxford University Press.