

Hard Determinism, Humeanism, and Virtue Ethics

Ben Vilhauer
William Paterson University

Abstract

Hard determinists hold that we never have alternative possibilities of action—that we only can do what we actually do. This means that if hard determinists accept the “ought implies can” principle, they must accept that it is never the case that we ought to do anything we do not do. In other words, they must reject the view that there can be “ought”-based moral reasons to do things we do not do. Hard determinists who wish to accommodate moral reasons to do things we do not do can instead appeal to Humean moral reasons that are based on desires to be virtuous. Moral reasons grounded on desires to be virtuous do not depend on our being able to act on those reasons in the way that “ought”-based moral reasons do.

Introduction

Hard determinists hold that determinism and moral responsibility are incompatible, that determinism is true for all our actions, and that we are therefore not morally responsible for any of our actions.¹ But most hard determinists do not wish to reject morality altogether. Most want a reformed version of morality that does not incorporate moral responsibility.² However, as part of their rationale for holding that determinism and moral responsibility are incompatible, hard determinists typically claim that determinism implies that we have no alternative possibilities of action. As they see it, if determinism is true, we cannot do anything but what we actually do. This means that hard determinists face a special task in making sense of moral reasons. In many ethical systems, agents’ most important moral reasons are based on “oughts”—they are, for example, propositions, beliefs, or facts about what agents ought to do.³ Such moral reasons will be referred to as deontic moral reasons in this

Ben Vilhauer is Assistant Professor at William Paterson University of New Jersey.

paper.⁴ It is widely agreed that “ought” implies “can.” If this is right, and if we cannot do things we do not actually do, then we cannot have deontic moral reasons to do things we do not actually do. This would imply that if hard determinists are to hold that we have moral reasons to do things we do not actually do, they require a nondeontic account of these reasons. The purpose of this paper is to provide such an account, in the form of a virtue-based, Humean approach to moral reasons.⁵ Humean moral reasons are based on desires to act virtuously. If an agent has a desire to act virtuously and believes that to act in some particular way would be to act virtuously, then that agent has a Humean moral reason to act in that way. It will not be claimed here that this approach to moral reasons is the only one that is consistent with hard determinism, just that it is one worth considering.

This paper is structured as follows. Section 1 argues that hard determinists risk a slide into compatibilism unless they accept that we never ought to do things we do not do and that it is therefore best for hard determinists not to include deontic moral reasons in their ethical theories. Section 2 argues that hard determinists can accommodate virtue and vice, even though they cannot accommodate praise and blame. Section 3 gives an initial argument for the claim that we can have Humean moral reasons to do things we do not do that are based on agents’ desires to act virtuously. Section 4 explores a modal difference between Humean and deontic moral reasons in more detail, in order to provide a more complete account of why deontic moral reasons are undermined by hard determinism but Humean moral reasons are not. Section 5 discusses the objection that desires to act virtuously can support only trivial moral reasons in cases where agents do not act virtuously. The main claim in this objection is that if an agent acts in some way that is not virtuous, then Humeans must accept that the agent’s strongest desire was not to act virtuously, so they must accept that the agent did not have most reason, all things considered, to act virtuously.

Sections 6 and 7 rebut this objection. According to Humeanism, agents need not do what they have the strongest desire to do, so an agent’s strongest desire may be to act virtuously even when that agent acts viciously. Further, agents need not have most reason, all things considered, to do what they have the strongest desire to do. This implies that even when agents’ strongest desire is to do something that is not virtuous, and they act viciously based on that desire, they may nonetheless have most reason, all things considered, to act virtuously. This will demonstrate that virtue-based, Humean moral reasons can be substantive even in cases where we do not act morally. The conclusion to be drawn is that a hard determinism based on this approach can satisfactorily explain how we have moral reasons to do things we do not do.

1. Hard Determinism and “Ought Implies Can”

The consensus in favor of the “ought implies can” principle seems to be as broad as any in ethics. “Oughts” are commonly understood as imperatives, or “commands of reason,” following Kant, and it seems nonsensical to command someone to do something he cannot do.⁶ If we accept hard determinism, it seems that agents cannot do anything but what they actually do. We never have any alternatives to the things we do because we are causally necessitated to do the things we do. On this view, the murderer cannot refrain from the murder. If hard determinists accept the “ought implies can” principle, they seem forced to accept that it is not the case that the murderer ought to have refrained from the murder. More generally, hard determinists seem forced to accept that agents never ought to do anything they do not do.⁷

Hard determinists might balk at accepting this, but they wind up in a difficult position if they do so. To see why, consider two strategies hard determinists might use to argue that agents sometimes ought to do things they do not do. One is based on rejecting the “ought implies can” principle, and another is based on accepting it. Both strategies have the disadvantage of pushing hard determinism close to the edge of a slippery slope toward compatibilism.

If hard determinists reject the “ought implies can” principle, then they can hold that agents sometimes ought to do things they do not do without inconsistency.⁸ That is, if “ought” does not imply “can,” then even if hard determinism implies that we can only do what we actually do, the fact that an agent does not do something does not imply that is not the case that he ought to do it. But “Agent A ought to do x ” is commonly taken to entail “If Agent A (knowingly) fails to do x , she is morally responsible for that failure.” Hard determinists can reject this entailment, but only if they do considerable work to explain why, if the claim that an agent could not have done something is no obstacle to the claim that the agent ought to have done that thing, it is nonetheless an obstacle to the claim that the agent is morally responsible for failing to do that thing. Even if such an explanation can be worked out, and the slide into compatibilism can be avoided, hard determinists might do well to consider easier paths.

Hard determinists who hold that agents sometimes ought to do things they do not do, but who accept the “ought implies can” principle, must also contend with the commonly accepted entailment just mentioned. But they face an additional problem all their own. They must seek to accommodate within hard determinism some sense of “can” sufficient to ground claims about what agents ought to do, and it is not clear that this can

be done. For example, they might take a page from the compatibilists' playbook, and adopt some version of the "conditional analysis" of "can," according to which an agent can do some action x so long as it is the case that if the agent chooses to do x , then the agent will do x .⁹ If "can" is understood in terms of the conditional analysis, then determinism need not conflict with the claim that an agent could have acted differently than she actually acted.¹⁰ The problem for this strategy is that taking this page from the compatibilists' playbook seems to force hard determinists to take the whole book. That is, if hard determinists accept a compatibilist analysis of the "can" in the "ought implies can" principle, then it is difficult for them to explain why the alternative possibilities that ground this "can" do not also satisfy the freedom requirement for moral responsibility. Hard determinists might just stipulate that there are two distinct senses of "can" at issue: one that is relevant for the "ought implies can" principle and another that is relevant for the freedom requirement of moral responsibility. But surely Occam's Razor cuts against this doubling of senses: theories are more attractive when adventitious senses are pruned away.

The arguments just made may not be enough for a definitive rejection of either of these strategies for preserving the claim that agents sometimes ought to do things they do not do. But they are enough to show that there are advantages available to hard determinists who give up this claim and, instead, provide a nondeontic account of agents' moral reasons to do things they do not do. The Humean, virtue-based account to be given here is one such approach. It bases moral reasons on agents' desires to act virtuously, instead of on claims about what they ought to do.

There may be other nondeontic approaches to moral reasons that would also be compatible with hard determinism. It may, for example, be possible to have a nondeontic account of moral reasons that invokes values rather than desires. That is, if there is such a thing as intrinsic value, it might be supposed that agents have moral reasons to act in ways that would actualize intrinsically valuable states of affairs, even if it is not the case that they ought to do so and they do not desire to do so. But this approach will not be considered here. It requires an account of how intrinsic value constrains moral reasoning, so to speak, which appeals neither to "oughts" nor to desires, and it is not obvious how this would work. That is, we might suppose that intrinsic value is such that agents ought to act in ways that actualize intrinsically valuable states of affairs. Alternatively, we might suppose that agents have reasons to actualize intrinsically valuable states of affairs insofar as they desire those states of affairs.¹¹ It may be that there is a third option, according to which intrinsic value gives agents moral reasons in a nondeontic way that does not depend upon what desires

agents happen to have.¹² But it is beyond the scope of this paper to explore such an option.

2. Hard Determinism and Virtue

The idea of a hard deterministic virtue ethics was first proposed by Spinoza.¹³ A hard deterministic virtue ethics claims that there is a basic distinction between virtuousness and viciousness, on the one hand, and praiseworthiness and blameworthiness, on the other. It also claims that hard determinists can ascribe the former characteristics to agents, though not the latter.

Some may object that it is impossible to ascribe vices or virtues to agents without blaming or praising them. But to hold that this is possible is just to distinguish between agents' possession of character traits, on the one hand, and agents being morally responsible for possessing those character traits, on the other. Perhaps the most intuitive starting point for grasping this distinction is to consider cases where someone's character is warped by childhood abuse. Someone may be caused to become cruel by childhood abuse or neglect; in such cases, it is clearly mistaken to blame him for his cruelty.¹⁴ According to a hard deterministic virtue ethics, all vices and virtues have etiologies that make all praise and blame similarly misplaced. To take a slightly more complicated example, suppose Doug was brought up by thieves, and as a result is deterministically necessitated to have the vice of thievishness, and to steal on some occasions. For hard determinists, the fact that Doug was deterministically necessitated to have this vice, and to act on it in some cases, means that he is not morally responsible for these things, and he therefore cannot be blamed for them. But this makes him no less vicious. There are many ways in which hard determinists can respond to Doug's viciousness without blaming him. For example, they can judge Doug's thievishness to be a vice, point his vice out to others, and watch their wallets when interacting with him. They can also exhort him to stop stealing and explain how stealing affects the people robbed, in the hope of causing him to change his attitude toward stealing and become more virtuous. None of these actions entail blame. Even punishing Doug does not imply blaming him, since there are teleological justifications of punishment that need not assume that the criminal to be punished is morally responsible for his crime (though any form of punishment that fair-minded hard determinists could accept would have to be demonstrably aimed at something like education). What is problematic about blame for hard determinists is the retributive dimension of blame, that is, the idea that Doug deserves to suffer because he has done vicious things. Hard determinists can accommodate nonretributive substitutes for our blaming practices that have

the goal of reducing the harm caused by vicious agents, or giving vicious agents reasons to become virtuous.

But how does such a virtue ethics allow us to do without “oughts”? On this point, we turn to Hume.

3. Hard Determinism and Humeanism

Everyone agrees that, according to Hume, reason does not tell us that we ought to pursue any ends.¹⁵ There is less agreement about whether Hume thinks instrumental rationality involves “oughts,” that is, whether he thinks we ought to take the means to our ends. The Humean hard determinism to be described in this paper is based on what might be called the “austere” interpretation of Humean practical rationality, according to which reason does not tell us that we ought to take the means to our ends.¹⁶ That is, according to austere Humeanism, there are no “oughts” at all in practical reasoning. According to this view, practical reasoning is just a combination of theoretical reasoning about what is the case, and desires.

A Humean account of moral reasons can be developed according to which moral reasons are practical reasons which are based on desires to act virtuously.¹⁷ We can explain how a desire to act virtuously yields reasons for action in an uncontroversial way using the widely accepted belief-desire model, which Hume implicitly advocates. That is, if some agent desires x , and believes that y is a means to x , then that agent has a reason to y .¹⁸ If Doug wants to act virtuously, and he believes that stealing is incompatible with acting virtuously, this gives him a reason not to steal.

At this point, let us consider and respond to two initial objections that might be made against this approach. First, it may be objected that supposing that moral reasons are based on desires to act virtuously, rather than on “oughts,” only pushes the problem back, because we cannot explain what it is to act virtuously without reference to the notion that virtuous people do what they ought to do. But this objection is mistaken because one need not use any deontic concepts to explain what it is to act virtuously. It is enough to describe the actions of virtuous people.¹⁹ One could then explain what it means to want to act virtuously as wanting to be such that one’s actions satisfy that description. There is no need to refer to any “oughts” in that description. Suppose virtuous people do not act cruelly. A deontologist might claim that virtuous people follow the principle that one ought not be cruel. A Humean could simply state that virtuous people do not act cruelly. There may be room for controversy about which people are the virtuous ones, but no more room than deontologists would have to allow for controversy about which actions are the ones people ought to take.

Second, it may be objected that the view that moral reasons are based on a desire to act virtuously implies that, when one acts on moral reasons, one's goal is morally second-rate in some sense. For example, it may be said that someone is more virtuous if she *simply wants to volunteer at a homeless shelter* than if she *wants to volunteer at a homeless shelter because it is virtuous*, and according to the theory presented here, only the latter desire can be the basis of a moral reason. The claim that one is more virtuous when one acts virtuously without aiming at virtue can be contested. But even if it is accepted, it need not be taken to be a problem for this theory. That is, it may be that perfectly virtuous agents have no moral reasons. They may have such well-adjusted desires that they never want to do anything that is not virtuous. The fact that the things that they want to do are virtuous may not be part of the content of their desires, and it may play no role in motivating them to do those things. In such cases, it seems strained to say that the reasons they have for their actions are in any distinctive sense moral reasons. It may be better to say (despite the risk of sounding romantic) that they have transcended moral reasons. It may be that only agents who are imperfectly virtuous, but who have desires that orient them toward greater virtue, can properly be said to have moral reasons.

4. The Modal Difference Between Deontic and Humean Moral Reasons

There is an important modal difference between deontic and Humean moral reasons, which can be explained as follows. If "ought" implies "can," then a deontic moral reason can only exist if it is possible for one to do what it is claimed one ought to do. So if hard determinism implies that we can only do what we actually do, and if "ought" implies "can," then we can have no deontic moral reasons to do anything we do not do. By contrast, the existence of a Humean moral reason depends only upon the existence of the desire to act virtuously and the belief that some particular action is a means to that end. This implies that Humean moral reasons are not undermined by hard determinism in the way that deontic moral reasons are because the existence of reasons grounded on desires does not depend on our being able to do what we have reason to do in the way that the existence of deontic moral reasons does.

Proponents of deontic moral reasons might object that what is important in thinking about determinism and deontic moral reasons to do things we do not do is not whether such reasons *exist*, but whether the agents who are *doing* the moral reasoning can consistently *believe* they exist. It might be argued that so long as the agents doing the moral reasoning do not know that

they will not act in some way, it is consistent for them to believe that they ought to act in that way.²⁰

This objection can be developed in more detail by considering the following scenario. Suppose that we are hard determinists who accept the “ought implies can” principle. Further, suppose that we know determinism is true, but we are unable to predict our future actions because the world is just too complicated for us to measure all the variables we would need to measure in order to make the necessary predictions. So we do not know what we will do, exactly, but we know that whatever it is, we are determined to do it. (This would not seem to be a metaphysically or epistemologically absurd supposition.) Suppose that Doug shares our epistemological predicament. Finally, suppose Doug is tempted to steal someone’s car, and suppose he is not sure whether he will yield to this desire or resist it.²¹

It appears that Doug can consistently believe that he has a deontic moral reason to refrain from stealing the car in this scenario, even if he is determined to steal the car. Let us consider how it looks from his point of view. He knows that if he steals the car, he is determined to steal it, so he cannot refrain from stealing it. If he cannot refrain from stealing it, then it is false that he ought to refrain from stealing it. But he does not know that he cannot refrain from stealing it, so he does not know that it is false that he ought to refrain from stealing it. It is therefore consistent for him to believe that he ought to refrain from stealing it.

Does this mean that hard determinists can accommodate deontic moral reasons to do things we do not do? Hard determinists may be able to use this point to accommodate such reasons in some formal sense, but their motivational efficacy would be undermined. If Doug cannot rule out the possibility that he cannot refrain from stealing the car, and he reasons consistently, he will have diminished confidence in his belief that he ought to refrain from stealing it. This has worrisome implications for Doug’s motivation. That is, human psychology would seem to be such that we are less strongly motivated to avoid acting in some way when we have diminished confidence in the belief that there is a moral reason to avoid acting in that way. If this is right (and if Doug is aware of no alternative, nondeontic moral reason to refrain from stealing the car) then Doug’s diminished confidence would tend to make him less motivated to resist the temptation than he would otherwise have been.

This argument does not require the claim that Doug can only have a deontic moral reason if he is *certain* that he can do what he believes he ought to do.²² It requires only the much weaker claim that the less confident one is in the truth of one’s belief that one has a deontic moral reason, the less motivationally efficacious that belief will tend to be. Consider Doug’s situation

as he stands beside the car, thinking about whether or not to steal it. Suppose he believes he ought not steal the car. If he accepts some theory of the will according to which he has alternative possibilities of action in making this choice, and he believes there are, for example, no hidden Frankfurtian controllers, then he would have no justification for the belief that he might not be able to refrain from stealing the car. In such a case, it would be unjustified for him to have diminished confidence in his belief that he ought not steal the car. But, as stipulated, Doug accepts hard determinism, so he *does* have a justification for the belief that he might not be able to refrain from stealing the car, and it *would* be justified for him to have diminished confidence in his belief that he ought not steal the car. If the claims made here are correct, this would tend to diminish the motivational efficacy of this belief.

Next, suppose that Doug succumbs to the temptation, hot-wires the car, and drives it away. How should he evaluate what he has done? He now knows that he could not have refrained from stealing the car, so he can conclude that it is not the case that he ought to have refrained from stealing it. In other words, in retrospect, Doug can know that there was no deontic moral reason not to act as he acted. The diminished confidence in the deontic moral reason that obtained in his prospective deliberations about how to act can be resolved in favor of knowledge that there was no deontic moral reason. This poses a problem for Doug's motivation to change his behavior in the future.²³ That is, it seems to be a fact about human psychology that we are more motivated to change the way we behave when we believe there *are* moral reasons not to behave as we have than when we believe there are *no* moral reasons not to behave as we have. If this is right (and if Doug is not aware of an alternative, nondeontic moral reason to refrain from stealing the car), then if he is tempted to steal a car again later, his belief that there was no deontic moral reason for him to refrain from the earlier theft would tend to make him less motivated to refrain from the later theft than he would otherwise have been.

This shows that, even if we consider things from the standpoint of epistemically limited agents like ourselves, the motivational efficacy of deontic moral reasons would be undermined by hard determinism. The motivational efficacy of Humean moral reasons is not similarly undermined. Consider Doug's situation again, but this time in terms of Humean moral reasons. Just as before, Doug is tempted to steal the car, and he is unsure whether he will steal it or not. He knows that if he steals it, he was determined to steal it. But it is clear that it would be vicious to steal it. The fact that this viciousness would be determined, and therefore blameless, does not make it any less vicious. If Doug wants to act virtuously, and he knows he has to refrain from stealing the car to act virtuously, then he has a Humean moral

reason to refrain from stealing it. When moral reasons are construed along Humean lines, the fact that Doug cannot rule out the possibility that he cannot refrain from stealing it does not throw the existence of the moral reason into question. Therefore, Doug's motivation to act morally need not be diminished by uncertainty about the existence of the moral reason, in the way that it is on the deontic account.²⁴

Suppose, again, that Doug succumbs to temptation, hotwires the car, and drives it away. According to the Humean account, what is the correct way for him to evaluate what he has done? The fact that his theft was determined makes it no less vicious, and since he wanted to act virtuously, he had a moral reason not to act as he acted. If Doug is tempted to steal another car later, the Humean account of his reason for his earlier action would not tend to undermine his motivation to act differently the second time in the way that the deontic account would.

5. Are Humean Hard Deterministic Moral Reasons Substantive?

Some may object that if Doug is deterministically necessitated to steal the car instead of acting virtuously, then he can only have a Humean moral reason in a trivial sense. This objection is partly based on a widespread concern about Humeanism that is not bound up in any special way with determinism. That is, according to the deontic view, agents' moral reasons need not be constrained by the desires they happen to have. (For example, on the Kantian view, moral reasons are based on the categorical imperative, and the categorical imperative is categorical precisely because it applies to agents irrespective of what they happen to desire.) For Humeans, by contrast, agents' desires are the only basis for reasons to act. It can seem natural to assume that, on the Humean view, whatever desire an agent actually acts on must be his strongest desire, and that acting on that desire must be what he has most reason to do, all things considered. If this is right, it follows that if Doug steals the car instead of acting virtuously, then his desire to steal the car was his strongest desire, and acting on it was what he had most reason to do. Since there is no other standard apart from his desires that can be used to assess what reasons he has, this might seem to imply that any reason he had for acting virtuously would have been, at best, a trivial reason—a reason that it would have been irrational for him to act on, all things considered.²⁵

One may have the concern about Humeanism just described independent of the worries about hard determinism at issue in this paper. But it may seem that the concern becomes even sharper in the context of hard determinism. For example, suppose one accepts a (highly speculative) metaphysics according to which quantum indeterminacy affects our deliberations about

what to do, thereby creating alternative possibilities of choice and action.²⁶ In that context, we might suppose that Doug has a desire to act virtuously, and a desire to steal the car, and that up until the moment when he decides to steal the car, it is metaphysically open which desire will actually motivate him to act. (Suppose that the outcome turns on whether a particular synapse fires at a particular time, and suppose that whether it fires at that time is sensitive to quantum indeterminacy.)²⁷ Under these assumptions, Doug would have Humean reasons both to steal the car and to act virtuously, and (at least prior to his decision to steal the car) there need be no fact of the matter about which desire is strongest, or which desire he has most reason to act on. So (at least prior to his decision to steal the car) this charge of triviality could not be leveled against his reason to act virtuously. Things might seem to be quite different if we accept hard determinism. In that case, if Doug steals the car, it is at no point metaphysically open which desire will actually motivate him to act. Even prior to his decision to steal the car, there can be clear facts of the matter about which of his desires are strongest, and which of his desires he has most reason to act on, all things considered, even if Doug does not know these facts.

It will be argued, however, that it does not follow from the fact that Doug steals the car that his desire to steal the car is his strongest desire, or that stealing the car is what he has most reason to do, and that the charge of triviality can therefore be blocked. But demonstrating this requires that the conception of practical rationality that underlies Humean hard determinism be set out in a bit more detail. The two elements of this conception that are needed to show why these conclusions do not follow are what will be called “motivational rationality,” on the one hand, and strength of will, or “volitional rationality,” on the other. Motivational and volitional rationality will be explained in sections 6 and 7, respectively.

6. Must Agents Do What They Most Strongly Desire to Do?

As mentioned earlier, according to the austere interpretation of Humeanism at work in the present paper, reason prescribes no “oughts” to govern conduct: the role of reason in practical rationality is limited to theoretical rationality. To be theoretically rational is to be able to make sound inferences, apply concepts correctly, and form consistent beliefs about causal relations. But theoretical reason has no bearing at all on one’s actions unless one happens to be what might be called a “motivationally rational” agent. If one is a motivationally rational agent, one’s motivation “tracks” the causal relations discerned by theoretical reason—one’s motivation with respect to the means to one’s

ends is responsive to one's beliefs about those means and one's beliefs about one's power over those means. Means are events that would be sufficient to cause (or constitute) ends.²⁸ If one has power over means to an end, then one has the ability to cause those means and, by extension, the ability to cause (or constitute) the end.²⁹ Using these terms, we can explain the sort of tracking involved in motivational rationality as follows: if one has an end, and has no conflicting end that one wants to achieve more, and one uses theoretical reasoning to form the belief that some event is the best means to that end over which one has power,³⁰ then one's motivation tracks that belief in the sense that one becomes motivated to cause that event. We might represent this in three steps:

- (1) The agent comes to desire end x more than any conflicting end.
- (2) The agent forms the belief that y is the best means to x over which she has power.
- (3) The agent becomes motivated to do y .

An agent who was theoretically rational but not motivationally rational with respect to end x could make it to (2) but not to (3).

It is entirely contingent from the standpoint of reason whether any agent is a motivationally rational agent, that is, there is no principle of reason according to which agents ought to be motivationally rational. Nonetheless, it is necessary that if some agent is motivationally rational, then she is motivated to take means y to end x (so long as she satisfies the conditions of not having some other end that conflicts with x and that she wants to achieve more, and forming the belief that y is the best means to x under her power). This is necessary in the sense that if we describe her as a motivationally rational agent and she is not motivated to take means y , despite satisfying the above conditions, then we are incorrectly applying the concept motivationally rational agent to her, and correctly applying concepts is a matter of theoretical rationality.³¹ If we stipulate that some agent has an end but is not motivated to take the means to it, despite satisfying the above conditions, then it is just an error of theoretical rationality to describe the agent as motivationally rational with respect to that end.

This account of motivational rationality makes it possible to reject the supposition that whatever means the agent is actually motivated to take is necessarily the means to the end that the agent most strongly desires to achieve. The fact that an agent does not take the means to some end when the conditions described above are satisfied can either mean that that end is not the one the agent most wants to achieve, or that the agent is not motivationally rational with respect to that end. For

example, suppose Doug is motivationally irrational because he is compelled by kleptomania to steal. His desire to achieve the end of acting virtuously might be stronger than his desire to achieve the end of stealing the car, but kleptomania might drive him to hotwire the car rather than walk away.

One cannot accept the possibility of such motivational irrationality if one presupposes a criterion for determining the strength of desires according to which the strength of a desire must be manifested in its effect on actual behavior. But there is little to recommend such a narrow criterion. If we accept an even slightly broader criterion, according to which the strength of a desire is instead manifested in its effect on possible behavior, then it becomes clear that we need not suppose that whatever means one is actually motivated to take is necessarily the means to the end one most strongly desires to achieve. We can explain a desire to achieve an end as a state of the agent that disposes the agent to be motivated to take the means to that end, and we can explain what it means for the agent to be disposed to be motivated to take the means to an end by describing a range of states of affairs in which the desire plays a role (along with beliefs) in causing the agent to be motivated to take the means to that end.³² If the agent's desire to achieve that end does not actually play a role in causing the agent to take the means to that end, then all the states of affairs in which it does play such a role are counterfactual states of affairs. According to hard determinism, such counterfactual states of affairs are physically impossible—that is, all counterfactual states of affairs are such that there was never a point in time at which it was physically possible for them to be instantiated in the actual world. But if the laws of nature in those counterfactual states of affairs are the same as those of the actual world, and the properties of the desires at issue are the same, then the counterfactual status of these states of affairs makes them no less explanatorily useful.

This counterfactual approach allows precise descriptions of the relative strengths of an agent's various desires to achieve ends because when describing the relative strengths of desires in the actual world, we can refer to their effects on behavior in counterfactual states of affairs. Consider again the claim that an agent desires to achieve end A more than he desires to achieve end B, but he is actually motivated to take the means to end B rather than the means to end A. With this counterfactual approach, we can explain that the desire to achieve end A is stronger in the following sense: in the actual world, some state of the agent is affecting the agent's will and "blocking" the motivational efficacy of his desire for end A, but in counterfactual states of affairs, which differ only in that this "blocking" phenomenon is absent, the agent is motivated to take the means to end A rather than the means to end B.³³

Once again let us consider Doug. Doug's desire to act virtuously is stronger than his desire to steal the car, but he hotwires the car instead of walking away because the motivational efficacy of his desire to act virtuously is "blocked" by his kleptomania (and he is therefore motivationally irrational). According to the counterfactual approach, what this means is that in counterfactual states of affairs in which the laws of nature are the same, and Doug is psychologically the same except that he is not afflicted with kleptomania, he walks away instead of hotwiring the car. This demonstrates that the means the agent is actually motivated to take is not necessarily the means to the end that the agent most strongly desires to achieve.

It might be objected that this counterfactual approach to determining the strength of desires is subject to a version of the "slippery slope toward compatibilism" argument made in section 1 against approaches to hard determinism that seek to preserve the claim that agents sometimes ought to do things they do not do. That argument went roughly as follows: if agents sometimes ought to do things they do not do, then why aren't they sometimes morally responsible for failing to do things they do not do? A similar argument might be posed against the counterfactual approach to determining the strength of desires, as follows: if physically impossible counterfactual states of affairs are relevant for determining what agents' strongest desires are, then why aren't they relevant for determining what agents can do (and by extension, for determining what agents ought to do, and are morally responsible for failing to do)?

This objection can be rebutted by emphasizing the distinction between dispositions and abilities. Hard determinists can understand some desires as dispositions to act in physically impossible ways without supposing that these dispositions ground abilities to act in physically impossible ways. The concept of disposition is one that lends itself naturally to analysis in conditional terms (i.e., "to have a disposition to do x is for it to be the case that if y had obtained, then one would have done x "), but it is a matter of great controversy whether the kind of abilities at issue in free will theory can be similarly analyzed.³⁴ According to hard determinism of the sort explained here, the kind of abilities at issue in free will theory cannot be similarly analyzed.

The purpose of this section has been to argue that Humean hard determinists can hold that agents need not do what they most strongly desire to do. In the situation considered here, Doug's strongest desire is to act virtuously, even though he steals the car, and he is motivationally irrational in failing to act virtuously. In making this argument, an important step has been taken toward showing that Humean moral reasons can be substantive moral reasons. One implication of this argument is that even if we accept the claim that we always have most

reason, all things considered, to do what we most strongly desire to do, we can sometimes have most reason to act virtuously even when we fail to act virtuously. But in the next section, it will be argued that we need not have most reason, all things considered, to do what we most strongly desire to do. This will make room for even further sophistication in Humean hard deterministic moral psychology.

7. Must Agents Have Most Reason to Do What They Have the Strongest Desire to Do?

In this section, it is argued that Humeans do not have to accept the claim that whatever agents most strongly desire to do is what they have most reason to do, all things considered. At the outset, it must be acknowledged that the Humean view cannot provide an account of what agents have most reason to do, all things considered, which correlates at all closely with the deontic view of this idea. As discussed above, this is because, on the Humean view, all practical reasons are based on agents' desires, while this is not the case on (e.g.) the Kantian view. The goal in the present paper is merely to give an account of what agents have most reason to do, all things considered, according to which this is not necessarily determined by what agents most strongly desire to do, but that nonetheless can be accommodated within Humean hard determinism.

Such an account can be developed by focusing on the phenomenon of weakness of will, that is, the phenomenon of failing to take the means to the end that one wills oneself to pursue (so to speak). Roughly put, the claim will be that Humean hard determinists can hold that what we have most reason to do, all things considered, is what we will ourselves to do, and in cases where the will is weak, one fails to do what one has most reason to do. Some work is required to show how Humeans can accommodate weakness of will, however, since the most common strategy for explaining weakness of will involves deontic beliefs. That is, weak-willed agents are commonly described as being torn between one end that they believe they ought to pursue and a conflicting end that they desire to pursue and that they actually pursue. Humeans cannot accept this description.

If, however, there is a way to interpret weakness of will as a conflict between desires, then it can be accommodated within Humeanism. Harry Frankfurt's distinction between first- and second-order desires and volitions offers an attractive approach to such an interpretation.³⁵ A second-order desire is a desire about a first-order desire, and a first-order desire is a desire about something other than a desire. A second-order volition is the kind of second-order desire one has when one wants a particular first-order desire to be one's will, that is, when one

wants a particular first-order desire to effectively motivate one to act.³⁶ According to the use to be made of Frankfurt's distinction in the present paper, an agent is weak-willed if the first-order desire which is the content of the agent's second-order volition is not effective in motivating the agent to act.³⁷ We might refer to the kind of lapse in practical rationality which is constituted by weakness of will as "volitional irrationality," to distinguish it from the motivational irrationality discussed earlier.

Second-order volitions can be used to give an account of what we have most reason to do, all things considered. Second-order volitions are desires that express our reflective evaluations of our desires. In other words, they express our judgments about what desires we desire to act on, all things considered. As Frankfurt puts it, when an agent forms a second-order volition, he "identifies himself" with one of his first-order desires.³⁸

Since desires ground reasons in the Humean theory, the process of forming a second-order volition can be seen as the process of reflecting on the various reasons for action we have and determining which of those reasons we want to be effective in motivating us to act. Also, since that higher-order volition is itself a kind of desire, it also grounds a reason. In other words, when we have a higher-order volition, we have a higher-order reason to act on a first-order reason. The point here is that Frankfurt's hierarchical picture of desires makes it possible to explain a sense of "most reason, all things considered," which is based not on the strength of desires but on a hierarchical endorsement of, and identification with, first-order desires.³⁹

Here is an example. Doug₁ has conflicting first-order desires. On the one hand, he has a first-order desire to act virtuously. On the other hand, he has a first-order desire to steal the car. He also has a second-order volition to be motivated by his first-order desire to act virtuously. But he is nonetheless not motivated to take the means to the end of acting virtuously: his first-order desire to steal the car is stronger than even the combination of his first-order desire to act virtuously and his second-order volition to act on his first-order desire to act virtuously. So he hotwires the car instead of walking away.

It should be noted that, since Doug₂ is motivated to take the means to the end he desires most, he is motivationally rational in stealing the car. Unlike our earlier kleptomaniacal Doug, Doug₂ is not motivationally "blocked," so there are no counterfactual states of affairs in which the laws of nature are the same, and where the strengths of his desires are the same, but where he walks away instead of hotwiring the car. The only counterfactual states of affairs in which the laws of nature are the same but Doug₂ is motivated to take the means to the end of acting virtuously are counterfactual states of affairs in which the combination of his first-order desire to act virtuously and

his second-order volition to be motivated by that first-order desire is stronger than it is in the actual world. Nonetheless, according to the Frankfurtian account just set out, Doug₂ has most reason, all things considered, to act virtuously, so he is volitionally irrational in stealing the car. The possibility of agents like Doug₂ demonstrates that agents need not have most reason, all things considered, to do what they have the strongest desire to do.

Volitional irrationality does not imply motivational rationality, however. Agents can be simultaneously motivationally and volitionally irrational. We can imagine such an agent if we return to our earlier kleptomaniacal Doug and give him a second-order desire to be motivated by his first-order desire to act virtuously. That is, suppose Doug has a first-order desire to act virtuously, a conflicting first-order desire to steal the car, and a second-order volition to be motivated by his first-order desire to act virtuously. Unlike Doug₂, Doug's first-order desire to act virtuously is stronger than his first-order desire to steal the car, yet he still hotwires the car instead of walking away because he is afflicted with kleptomania. In counterfactual states of affairs in which the laws of nature are the same, and Doug is psychologically the same except that he is not afflicted with kleptomania, he walks away instead of hotwiring the car.

Conclusion

The goal in this paper has been to demonstrate that a virtue-based, Humean account of moral reasons allows hard determinists to explain how we can have moral reasons to do things we do not do. If hard determinism implies that we only can do what we actually do, and "ought" implies "can," then hard determinism implies that we have no deontic reasons to do anything we do not actually do. But hard determinism does not similarly undermine Humean moral reasons because the fact that one desires to act virtuously does not imply that one can act virtuously. Humean hard determinism can support a rich moral psychology. It can explain not only how we can have moral reasons to do things we do not do but also why it isn't necessary that agents act on their strongest desires, and how we can have most reason, all things considered, to do something other than what we most strongly desire to do. This demonstrates that virtue-based, Humean moral reasons can be substantive moral reasons.

Notes

Earlier versions of this paper were presented at the Central States Philosophical Association and the Southern Society for Philosophy and Psychology (both in 2006) and at the New Jersey Regional Philosophy Association in 2005. Thanks to Tomis Kapitan, John Peterman, Pete

Mandik, Eric Steinhart, Barbara Andrew, Amy Kind, Alex Rajczi, and the anonymous reviewers at *The Southern Journal of Philosophy* for helpful comments.

¹ Hard determinists need not deny that there are indeterministic events at the quantum level. They need only deny that we have good reason to believe that quantum indeterminacy propagates upward to the macrophysical level in ways that have significant effects on human actions.

² To hold people morally responsible is to praise or blame them. But praise and blame are not the whole of morality, so it seems at least *prima facie* possible to reject moral responsibility without rejecting the whole of morality. Examples of other work that argues that this is possible include the present author's "Hard Determinism, Remorse, and Virtue Ethics," *Southern Journal of Philosophy* 42, no. 4 (2004): 547–64; Michael Slote's "Ethics Without Free Will," *Social Theory and Practice* 16, no. 3 (1990): 369–83; Bruce Waller's *Freedom Without Responsibility* (Philadelphia: Temple University Press, 1990); and Derk Pereboom's *Living Without Free Will* (Cambridge: Cambridge University Press, 2001). (Pereboom describes himself as a "hard incompatibilist" rather than as a hard determinist—he does not claim that determinism is true, but he claims that moral responsibility is incompatible with both determinism and the sort of indeterminism that would result if quantum indeterminacy had significant effects on actions, and he thinks we have good reason to believe that one or the other of these obtains.) For such philosophers, the crucial point of disagreement with compatibilists (who hold that moral responsibility is compatible with determinism) is about praise and blame, not about morality as such.

³ See Jim Pryor, "Reasons and That-Clauses," *Philosophical Issues* 17 (2007): 217–44, for a helpful discussion of possible views on the ontology of reasons.

⁴ The term "deontic moral reasons" is used in a variety of ways in contemporary ethics that differ from the way it is used here. For example, this term is sometimes used to distinguish obligatory reasons from supererogatory reasons (the former being referred to as deontic moral reasons and the latter being referred to as nondeontic moral reasons). But some ethicists regard obligatory reasons and supererogatory reasons as two kinds of "ought"-based reasons, the former being stronger than the latter. Understood this way, both obligatory and supererogatory reasons would count as deontic moral reasons as the term is used in this paper. (Some might argue that if supererogatory reasons are "ought" based, then the "ought implies can" principle fails for them, but this issue is beyond the scope of the present paper.)

⁵ The approach to hard determinism to be considered here does not incorporate Hume's own theory of free will. Hume himself was a compatibilist, not a hard determinist. Not much time will be spent in this paper arguing against Hume on this point, but here is a brief description of his position and an objection to it. In *An Enquiry Concerning Human Understanding* (ed. L. A. Selby-Bigge [Oxford: Oxford University Press, 1975]), Hume seems to have two main reasons for thinking that determinism is compatible with moral responsibility. One is his compatibilist account of the ability to do otherwise (Section 8, Part 1). Hume defines free will ("liberty," in his terms) as follows: "By liberty ... we can only mean a *power of acting or not acting*

according to the determinations of the will; that is, if we choose to remain at rest, we may; if we choose to move, we also may." In context, it is clear that, though Hume is endorsing the idea that we might have chosen differently, and that if we had chosen differently, we would have acted differently, he is not endorsing the idea that we have alternative possibilities of action at any point in time in the actual world. Hume's idea is if something had been different about the past or the laws of nature, then we would have been causally necessitated to choose differently. Hume's view seems to be that a killer could have chosen not to kill just in case he does choose not to kill in some other possible world with the right sort of relationship to the actual world. This account of "can" anticipates some contemporary compatibilist accounts (see the main text below and note 10), but it is not what we pretheoretically mean by "can" when we condition a claim that the killer is morally responsible for the killing on the claim that he could have refrained from the killing. In this context, we mean that the killer could have chosen not to kill in the particular situation in the actual world in which he did in fact do the killing, not that he chooses not to kill in some other possible situation. To make the point more concrete, the point is that if we justify blame by claiming that the killer could have refrained from stabbing his victim, we are claiming that he could have refrained from making these particular, actual stab wounds in the victim, not that the killer refrains from making some other stab wounds in some other possible world. Hume's second main reason for holding that determinism is compatible with moral responsibility is his view that, as a matter of empirical fact, we would still hold unethical agents morally responsible even if we believed that determinism was true (Section 8, Part 2). This is problematic because the views just discussed commit Hume to the position that there is a substantive metaphysical question about whether free will is compatible with determinism, which philosophers can answer in the affirmative only if they provide an acceptable account of the modal structure of alternative possibilities of action which makes them compatible with determinism. This implies that if we did not have such an account, but we continued to hold people morally responsible despite believing in determinism, we would be making an error.

⁶ It may well be that "x ought to do y" implies both (A) that x can do y and (B) that it is possible for x not to do y. That is, if it is incoherent to command x to do something he cannot do, it may be equally incoherent to command x to do something that he will inevitably do. (Kant appears to express this view at *Groundwork* Ak. 414.) It is not clear, however, that the broad consensus supporting A also exists for B. So only A will be addressed in the main text.

⁷ Many free will theorists agree that hard determinism entails that it is not the case that agents ought to do things they do not to do. See, e.g., I. Haji, "Moral Anchors and Control," *Canadian Journal of Philosophy* 29 (1999): 175–203. Saul Smilansky rejects this view, however: he divides what he calls "common sense morality" into "two distinct parts": "substantive morality," on the one hand, which concerns "what morally ought to be done (or not done)," and the "accountancy" part of morality, on the other, which concerns the positive reactions people ought to get when they do what ought to be done, and the negative reactions people ought to get when they do what ought not to be done. According to Smilansky, hard determinism undermines the

accountancy part of morality but not the substantive part of morality. But Smilansky does not explain how hard determinists who hold that agents sometimes ought to do things they do not do can handle the “ought implies can” principle. See his paper “The Ethical Advantages of Hard Determinism,” *Philosophy and Phenomenological Research* 54 (1994): 355–63 (quotes from pages 356–57).

⁸ John Martin Fischer, a compatibilist, is perhaps the most prominent proponent of rejecting the “ought implies can” principle in contemporary free will theory. See his papers “‘Ought-implies-can’: Causal Determinism and Moral Responsibility,” *Analysis* 63, no. 279 (2003): 244–50, and “Recent Work on Moral Responsibility,” *Ethics* 110 (1999): 93–139.

⁹ Thanks to Tomis Kapitan for suggesting this possibility.

¹⁰ Suppose that the agent is not subject to coercion or to any bodily constraint, so that her choices of how to act are efficacious in causing her to act. Under these conditions, even given determinism, it is true that, if the agent had chosen to act differently than she actually acted, then she would have acted differently than she actually acted. (The conditional analysis is a descendent of Hume’s own theory of free will, but that part of Hume’s theory is not part of the theory defended here. Also see note 5.)

¹¹ Hume himself of course rejects the idea of intrinsic value (see, e.g., *Treatise* 1.3.1). But one can include intrinsic value within one’s theory of moral reasons without giving up the Humean view that all reasons are based on desires. In such a theory, the idea of intrinsic value could be used to explain which desires virtuous agents have. That is, Humeans must insist that agents only have reasons to bring intrinsically valuable states of affairs into existence insofar as they desire such states of affairs, but they might nonetheless understand virtuous agents as the ones who have such desires. A Humean theory constructed along these lines would make intrinsic value explanatorily prior to virtue. But this is not the only way to construct a Humean theory. A Humean theory may (following Hume himself) entirely dispense with the idea of intrinsic value. It may also hold that the best way to understand the intrinsically valuable is in terms of the ends of virtuous agents, in such a way that virtue is explanatorily prior to intrinsic value. Also see note 18.

¹² It should be noted that this question is not the same as the question of whether there are supererogatory reasons. Some ethicists understand supererogatory reasons as weak, nonprescriptive “oughts,” and understood this way, they are a kind of “ought”-based moral reasons.

¹³ See Michael Slote’s discussion of Spinoza in “Ethics Without Free Will.” (Spinoza’s approach is egoistic in a way that contrasts with both the virtue ethics Slote advocates and the virtue ethics set out in the present paper.)

¹⁴ Slote offers a similar example (“Ethics Without Free Will,” 377).

¹⁵ See, e.g., *Treatise* 2.3.3, where Hume claims that it is not “contrary to reason to prefer the destruction of the whole world to the scratching of my finger.”

¹⁶ For recent discussions about whether Hume thought there were “oughts” of instrumental reason, see Christine Korsgaard, “The Normativity of Instrumental Reason,” in *Ethics and Practical Rationality*, ed. Garrett Cullity and Berys Gaut (New York: Clarendon, 1997) and “Skepticism about Practical Reason,” *Journal of Philosophy* 83, no. 1

(1986): 5–25. Also see Elijah Millgram, “Was Hume a Humean?” *Hume Studies* 21 (1995): 75–93.

¹⁷ Support for this reading of Hume can be found in a famous passage in *An Enquiry Concerning the Principles of Morals*, Appendix 1:

It appears evident, that the ultimate ends of human actions can never, in any case, be accounted for by *reason*, but recommend themselves entirely to the sentiments and affections of mankind, without any dependence on the intellectual faculties. Ask a man, *why he uses exercise*; he will answer, *because he desires to keep his health*. If you then enquire, *why he desires health*, he will readily reply, *because sickness is painful*. If you push your enquiries farther, and desire a reason, *why he hates pain*, it is impossible he can ever give any. This is an ultimate end, and is never referred to any other object.... It is impossible that there can be a progress; and that one thing can always be a reason, why another is desired. Something must be desirable on its own account, and because of its immediate accord or agreement with human sentiment and affection. Now virtue is an end, and is desirable on its own account, without fee or reward, merely, for the immediate satisfaction which it conveys[.]

¹⁸ It is common usage to refer to this sort of theory as Humean, but its prominence in contemporary philosophy owes as much to Donald Davidson as to Hume. See, e.g., Davidson, “Actions, Reasons, and Causes,” *Journal of Philosophy* 60, no. 23 (1963): 685–99.

¹⁹ If there is such a thing as intrinsic value, then a Humean account of virtuous action can be given that does not stop with such description but goes on to explain virtuous action as action aimed toward actualizing intrinsic value. Such an account will not be developed here, and it is not obvious that it would add anything substantial to our explanation because it might not give us any way of specifying virtuous actions that did not depend on references to virtuous agents. Also see note 11.

²⁰ Thanks to Tomis Kapitan for pointing out this objection.

²¹ It might be thought that hard determinism implies that Doug’s deliberations can have no effect on his actions because it is inevitable that he is going to act in a certain way, so there is no point in his deliberating about whether or not to resist the temptation. But that impression is mistaken. Suppose Doug is standing beside the car, reflecting on two options: to jimmy the door and hotwire the car, or to walk away. And suppose he reasons that it would be more virtuous to walk away, and he walks away. Hard determinism of course implies that it is inevitable that he will decide to walk away. But (barring concerns about epiphenomenalism that are entirely independent of hard determinism) his decision to walk away is nonetheless a part of the causal sequence that terminates in his walking away, just as much as the contractions of the muscles that move his legs are parts of the causal sequence. Hard determinism does not imply that some other process causes him to walk away and prevents his reasoning from having an effect. So moral reasoning can still be a cause of moral actions, and therefore it still has a point.

It might also be objected that the idea of temptation does not make sense in the context of hard determinism. But hard determinists can

accept that agents sometimes experience conflicts between desires they judge to be moral and desires they do not judge to be moral, and this is enough to make sense of the idea of temptation. If determinism is true, then the outcome of any such conflict is inevitable, but the agent in the midst of that conflict need not know what the outcome will be. (The discussion of volitional irrationality in section 7 can be read as a more detailed account of how hard determinists can understand temptation.)

²² Thanks are due to Tomis Kapitan for an objection which prompted this point.

²³ Changing one's behavior is not incompatible with hard determinism. There can be patterns in one's behavior that persist over intervals of time and then change, even if determinism is true.

²⁴ An objection might be made here as follows. In order to form an intention to do x , one must believe that it is possible for one to do x . If one holds that the motivational efficacy of a belief is proportionate to one's confidence in the truth of that belief, then it would be natural to make the parallel assumption that one's ability to form an intention to x is proportionate to one's confidence in the truth of one's belief that it is possible for one to do x . This suggests that, even for the Humean hard determinist, there may be an erosion of one's ability to act intentionally that parallels the erosion of motivational efficacy of beliefs about what one ought to do that afflicts the deontological hard determinist. That is, the Humean hard determinist must accept that, though he wants to act virtuously, he may not be able to act virtuously, and this may in some cases diminish his ability to intend to act virtuously. In response to this objection, consider the following. Both the Humean hard determinist and the deontological hard determinist must contend with any erosion of ability to intend to do x that is consequent upon recognizing that they may well be unable to do x . But the deontological hard determinist must also, simultaneously, contend with an erosion of the motivational efficacy of his moral reason to do x , i.e., of his belief that he ought to do x . Since the Humean hard determinist's moral reason to do x is not ought based, it is not subject to the same erosion. So the deontological hard determinist must contend with more sources of erosion of his capacity to act intentionally on moral reasons than the Humean hard determinist. (Thanks to Tomis Kapitan for a comment that prompted this point.)

²⁵ Hume himself might agree with this objection, but as noted before, no claim is made here that the present approach hews faithfully to Hume in all respects.

²⁶ See note 1.

²⁷ Robert Kane develops a metaphysics something like this in *The Significance of Free Will* (New York: Oxford, 1996), but not in the context of Humeanism about moral reasons.

²⁸ "Events" is intended here in the broad sense of "things that happen." In everyday usage, "means" often refers to things that we would pretheoretically think of as physical objects, i.e., things *to which* things happen, rather than things *that* happen. We might call a shovel (which it may be natural to represent as a physical object rather than an event) a means to the end of having a hole in the ground. But in such cases, the physical object is always of interest because of the role it can play in events that will bring the end into existence. The shovel is of interest because of the role it plays in shoveling. So from the

theoretical standpoint, it does not distort anything to regard the events these objects are used to cause as the means, rather than the physical object. On a different point, it is sometimes better to see means as constituents rather than causes of ends because the relation between means and ends is sometimes too immediate for it to make sense to see one as the cause of the other. Suppose one has the end of *getting some breakfast*, and the means one chooses is *ordering pancakes at the diner*. It seems better to say that ordering pancakes at the diner constitutes getting some breakfast than to say that it causes getting some breakfast.

²⁹ One gets beliefs about means by starting with one's beliefs about the state of affairs that is one's end, along with one's beliefs about the laws of nature, and drawing conclusions about various possible events that would suffice to cause one's end. One gets beliefs about one's power over means to one's ends in a similar way—i.e., one starts with one's beliefs about the means and one's own causal powers, along with one's beliefs about the laws of nature and the state of affairs in which one currently is, and draws conclusions about one's own power over the means.

³⁰ It is beyond the scope of the present paper to present a detailed account of what it is for a means to be the best means over which one has power. Efficiency will be an important consideration for most agents. It may be that virtuous agents seldom or never act in ways that use agents as mere means. Since these attitudes toward means can be understood as desires for ends (i.e., desires to act efficiently or virtuously), they can be incorporated into the Humean theory without difficulty. It is important to note, however, that hard determinism implies that if an agent had complete knowledge about the laws of nature and the state of affairs she was in, and she reasoned consistently in view of all this knowledge, she would conclude that there was never more than one means over which she had power at any given time. Since it is unlikely that human agents will ever have that kind of knowledge, however, it is likely that we will always deliberate about multiple means, all of which are under our power to the best of our knowledge.

³¹ Korsgaard appears to think that one can only hold that it is necessary for motivationally rational agents to be motivated to take the means to the ends they want most if one also holds that instrumental rationality has imperatival status, i.e., that practically rational agents ought to take the means to their ends. But she offers no argument for this. See Korsgaard, "The Normativity of Instrumental Reason."

³² This of course assumes the desire at issue is not a "finkish" disposition, i.e., a disposition that might be eliminated in exactly those circumstances that would ordinarily cause it to be manifested. See David Lewis, "Finkish Dispositions," *Philosophical Quarterly* 47 (1997): 143–58. Also see Kadri Vihvelin, "Free Will Demystified: A Dispositional Account," *Philosophical Topics* 32: 427–50.

³³ See note 30. Korsgaard calls these cases of "true irrationality." She lists "rage, passion, depression, distraction, grief, [and] physical or mental illness" as examples of psychological states that can block the motivational efficacy of desires to pursue ends. See Korsgaard, "Skepticism about Practical Reason," 12–13.

³⁴ See section 1 and notes 5 and 10.

³⁵ Harry Frankfurt, "Freedom of the Will and the Concept of a Person," *Journal of Philosophy* 68, no. 1 (1971). Frankfurt introduces this distinction in the service of a compatibilist account of free will, but accepting this distinction in no way entails compatibilism. Also see note 36.

³⁶ We can also follow Frankfurt in allowing that volitions may be formed at a higher order than the second. But this additional complexity will not be addressed in the main text.

³⁷ According to Frankfurt, an agent has free will if her second-order desire is effective in motivating her to act. However, Frankfurt's distinction between first- and second-order desires in no way entails his compatibilism.

³⁸ See, e.g., Frankfurt "Freedom of the Will and the Concept of a Person," 13.

³⁹ Even further sophistication can be built into this account without violating the desire-basing limitations of Humeanism. For example, one might instead hold that what one has most reason to do, all things considered, is determined by the second-order desire one would have if one had all the relevant information, thereby accommodating cases like Williams's gin-petrol case.