UTILITARIAN MORAL VIRTUE, ADMIRATION, AND LUCK

Robert J. Hartman

Published in *Philosophia* (2015) 43: 77-95

Abstract: Every tenable ethical theory must have an account of moral virtue and vice. Julia Driver has performed a great service for utilitarians by developing a utilitarian account of moral virtue that complements a broader act-based utilitarian ethical theory. In her view, a moral virtue is a psychological disposition that systematically produces good states of affairs in a particular possible world. My goal is to construct a more plausible version of Driver's account that nevertheless maintains its basic integrity. I aim to accomplish this goal by developing four problems concerning admiration and luck for Driver's account. Subsequently, I modify the account in a way that partially or entirely mitigates each difficulty. Finally, I attempt to undermine Driver's rationale for rejecting the recommended modification and explore how well the modified account of moral virtue fits with utilitarian accounts of right action.

Every tenable ethical theory must have an account of moral virtue and vice. Julia Driver has performed a great service for utilitarians by developing a utilitarian account of moral virtue that complements a broader act-based utilitarian ethical theory. In her view, a moral virtue is a psychological disposition that systematically produces good states of affairs in a particular possible world. My goal is to construct a more plausible version of Driver's account that nevertheless maintains its basic integrity. I aim to accomplish this goal by developing four problems concerning admiration and luck for Driver's account. Subsequently, I modify the account in a way that partially or entirely mitigates each difficulty. Finally, I attempt to undermine Driver's rationale for rejecting the recommended modification and explore how well the modified account of moral virtue fits with utilitarian accounts of right action.

In the first section, I explicate the relevant features of Driver's account of moral virtue. In the second and third sections, I marshal four pre-theoretical moral intuitions against Driver's account concerning the way in which moral virtues are worthy of admiration and the way in

which luck should not affect the moral status of dispositions.[1] After each argument, I consider how one sympathetic to Driver's account might respond on its behalf. My assessment, however, is that Driver's account lacks the internal resources required to undermine the total force of these problematic intuitions. Subsequently, I note that Aristotelian and Kantian conceptions of moral virtue are able to endorse those intuitions. As a result, a proponent of Driver's account cannot dismiss these arguments as problems for everyone.

In the fourth section, I argue that Driver's basic account of moral virtue becomes more plausible if it is altered in a particular way. My suggestion is to widen the scope of consequences relevant for determining the moral status of a disposition (whether it is a moral virtue, moral vice, or neither) from the consequences that it produces in a single possible world to those it produces in a possible world and worlds close to it.[2] I argue that the modified account either partially or entirely assuages each difficulty, and thus the modified account better captures the pre-theoretical idea of moral virtue than Driver's original account. Since one criterion for the plausibility of an account of moral virtue concerns how well the account captures the pre-theoretical idea, modifying Driver's account in this way increases its plausibility.

In the fifth section, I consider why Driver limits the consequences that count toward the moral status of a disposition to those that it produces in a single world. I argue that there is at least one way in which Driver's rationale for the restrictive limit in fact supports the

[1] The modest *pre-theoretic* status of the intuitions safeguards these arguments from begging the question against the utilitarian; for while I think that these intuitions survive reflection, claiming so without substantial argument would beg the question.

[2] Ben Bradley (2005) also recommends this modification. My contribution is three-fold. First, I advance several new arguments against Driver's view, and I show how the modified account either partially or entirely undermines each problem. Second, I provide new reasons for thinking that the permissive limit on consequences is more plausible than Driver's restrictive limit. Third, I also explore how well utilitarian accounts of right action fit with the modified account of moral virtue.

recommended permissive limit. In the sixth section, I explore how well the modified account fits with various utilitarian accounts of right action.

Driver's Account of Moral Virtue

Driver (2001, p. 82) characterizes a moral virtue as a "character trait that produces more good (in the actual world) than not systematically." Likewise, a moral vice is a character trait that systematically brings about more bad states of affairs than not in the actual world. There are four aspects of these compact definitions that require elaboration before I advance arguments against them. First, Driver (2001, pp. 95-106) identifies 'the good' with flourishing of beings in the moral community, and her account is not committed to any particular account of flourishing. Second, to be a moral virtue, a disposition must systematically produce good states of affairs. As such, moral virtues are causally and regularly related to good consequences.[3] This requirement allows moral virtues to bring about some bad consequences, and it also allows some moral virtues to be better than others (Driver 2001, p. 74). The better moral virtues are those dispositions that systematically produce more good than others.[4]

Third, moral virtue and vice are supervenient properties, and the supervenience base of a moral virtue is a disposition that systematically produces good states of affairs. In other words, moral virtue and vice respectively supervene on character traits that systematically produce good or bad consequences, because (i) the moral status of a character trait depends on the

[3] In response to an argument from Onora O'Neill (2004), Driver (2004, pp. 35-36) exchanges the causal requirement for a weaker 'make a difference' requirement. Since causation is one way to make a difference, I retain the original language of causation with the understanding that she is committed only to the thinner requirement.

[4] There is an ambiguity in the claim that the better virtues systematically produce more good than other virtues. The phrase 'produce more good' could mean either (i) produce more overall good in the world or (ii) produce a higher average quantity of good each time it is actualized. As Bradley (2005, pp. 296-297) recognizes, if we assume the meaning that the better virtue produces more overall good in the world, then we get the counterintuitive result that wit may be a better moral virtue than generosity only because it is far more frequently enacted. The more plausible disambiguation is that one moral virtue is better than another because it systematically produces a higher average quantity of good each time it is actualized.

consequences that the trait produces and (ii) there can be no change in the moral status of a character trait without some alteration in the consequences it produces. Fourth, there may be more than one supervenience base within a particular world but a world is the limit (Driver 2001, p. 82, 85). Thus, the consequences that character traits produce in other possible worlds—even close worlds—do not count toward the moral status of character traits in the actual world. In summary, the moral status of a disposition *x* in world *w* supervenes on the consequences that instances of *x* produce in *w*.

Over the next two sections, I argue that Driver's account has at least four counterintuitive ramifications. In some of those cases, restricting the supervenience base of moral virtues and vices to the consequences produced in only one world saliently generates the problem.

Arguments from Admirable Psychology

Driver acknowledges two ways in which the psychological structure of a disposition is relevant to its being a moral virtue or vice. First, a necessary condition for a disposition's being a moral virtue is that it is a psychological disposition (Driver 2001, p. 106). As a result, even if a non-psychological disposition such as smelling good or having dark hair systematically produces good states of affairs, it cannot be a moral virtue. Second, pre-theoretical good intentions (from now on, "good" intentions) are contingently related to a disposition's status as a moral virtue (Driver 2001, p. 56). Dispositions that characteristically give rise to "good" intentions reliably produce good states of affairs, because people are typically successful in bringing about the good for which they aim. Nevertheless, moral virtues do not necessarily give rise to "good" intentions; it is possible for a disposition that characteristically gives rise to pre-theoretical bad intentions (from now on, "bad" intentions) to be a moral virtue if that disposition systematically produces

good states of affairs (Driver 2001, p. 53). So then, while moral virtues are necessarily psychological dispositions, they are not necessarily dispositions involving any particular psychology such as "good" intentions or motivations.

In order to test the intuitive plausibility of the claim that moral virtues are not necessarily dispositions to form "good" intentions, Driver offers a thought experiment about a disposition to form "bad" intentions that satisfies her criteria of moral virtue. And because she (2001, p. 56) thinks that as a contingent matter of fact all actual human moral virtues are dispositions to form "good" intentions, the content of the thought experiment appeals to non-human persons.[5] Suppose there exists a society of beings called Mutors who evolved differently from humans (Driver 2001, pp. 55-56). There are two peculiar facts about them. First, the life expectancy of Mutor children is doubled if they are physically beaten when they are 5.57 years old. Second, some Mutors have the intense desire to beat Mutor children when and only when they are 5.57 years old. Call that small group of Mutors with the intense desire to beat children at the relevant age *the beaters*, and the trait they possess *the beater trait*. In this society, children are brought to the beaters when they are 5.57 years old in order to have them beaten, because it is very good to have one's lifespan doubled. The beaters, however, violently flog the children solely because they enjoy the activity. The good effect on the children is psychologically irrelevant to the beaters.[6] Since the good consequences of the beating systematically outweigh the bad, the beater trait is a moral virtue according to Driver's account.

Intuitively, however, no one should agree that the beater trait is a moral virtue, because this trait is a disposition to hurt children for the sheer pleasure of it. In fact, Driver (2004, p. 40)

[5] Since Driver (2001, p. 56) takes the concept of moral virtue to apply across species, the fact that her thought experiment appeals to non-human beings is unproblematic.

[6] We may assume that the beaters are ignorant of lifespan enhancing effects of the beatings.

shares the intuition that the beater trait is not a moral virtue, which is why she offers several arguments to discredit that intuition. But before we consider her counter-arguments, I want to explore some general explanations for why we have the intuition that the beater trait is not a moral virtue. One explanation is that the following principle captures a core feature of the pre-theoretical idea of moral virtue: necessarily, no moral virtue is a disposition to form "bad" intentions. Another explanation concerns a stronger principle: necessarily, moral virtues are dispositions of the mind and will worthy of admiration. The latter principle is stronger because it entails but is not entailed by the former. That is, a disposition's being worthy of admiration entails that it is not a disposition to form "bad" intentions. The converse, however, is false, because there are dispositions that do not regularly lead to the formation of "bad" intentions which are not worthy of admiration. One might think that the stronger principle most fully captures a core aspect of the pre-theoretical idea of moral virtue. This contention receives some inductive support from Driver's (2001, p. 56) claim that all moral virtues in the actual world are dispositions to form "good" intentions.[7] Even if the stronger principle is true, my argument depends only on the weaker claim that necessarily, no moral virtue is a disposition to form "bad" intentions. And because Driver's account of moral virtue is incompatible with the relevant case intuition and the weaker principle, her account cannot accommodate this core aspect of the pre-theoretical idea of moral virtue.

Consider another argument from admirable psychology against Driver's account. Suppose that the beater trait systematically produces an average of *x* goodness overall (where *x* is some positive value) each time the trait is activated. Suppose also that benevolence systematically produces an average of *x* goodness overall each time it is activated in that same

[7] In conversation, Driver confirmed this interpretation.

world. The beater and benevolent traits are thus equally instrumentally morally valuable. As a result, two agents who are identical but for the fact that one is benevolent and the other is a beater have equally instrumentally valuable sets of traits. It is, however, pre-theoretically plausible that benevolence is a morally better disposition than the beater trait in that world. For that very reason, it is also pre-theoretically plausible that the benevolent agent is a better person than the near identical agent with the beater trait. Consider a general principle that captures the relevant pre-theoretical intuition: necessarily, a disposition with psychology worthy of admiration is morally superior to a disposition with "bad" psychology for two equally instrumentally valuable dispositions in the same world.[8] Driver's account of moral virtue cannot accommodate this intuitive reaction or the general principle that captures it, because, according to her account, the moral value of a disposition is exhausted by the good it produces in a world. This implies that neither disposition nor agent is morally superior to its complement. Let us now turn to consider and assess a response to these arguments that a utilitarian might be tempted to advance as well as the responses that Driver offers.

One sympathetic to Driver's account might respond to these arguments by suggesting that the important feature of her account is that it is a distinctively utilitarian analysis of moral virtue, and, as such, the inability of her account to accommodate some of the pre-theoretical idea of moral virtue is unproblematic. It is important to see why Driver does not respond in this way. Her (2001, p. 62; 2004, p. 33) project is to provide a utilitarian account of moral virtue that is as good as or better than its Aristotelian and Kantian rivals. And one criterion for measuring the plausibility of various conceptions of moral virtue concerns how well each conception accounts

[8] W. D. Ross (1930, pp. 34-35) offers a similar argument regarding the morally significant features of an action. Should one abstain from lying and produce 1,000 units of happiness or tell a lie and produce 1,001 units of happiness? It is intuitive that the former act should be pursued, even though that action is slightly less instrumentally valuable, because this act has moral significance beyond its instrumental value.

for the pre-theoretical idea of moral virtue. Thus, the inability of Driver's account to capture core features of the pre-theoretical idea poses a problem for her goal of offering an account of moral virtue that is more plausible than other major conceptions.[9] As a result, she attempts to undermine the aspects of the pre-theoretical idea that her account cannot accommodate by offering an error theory and an analogical argument.

In her first response, Driver (2001, p. 56; 2004, p. 40) offers an error theory to undermine the intuition that the beater trait is not a moral virtue—which is an intuition that she admittedly shares.[10] According to the error theory, our intuitions about the Mutor Case are infected by our intuitions about human traits in two ways. First, people are incorrectly disposed to think that "good" intentions are necessary for moral virtue, because there is an extremely close but contingent correlation between "good" intentions and moral virtues in actual human beings. As a result, we are disposed to think that the beater trait cannot be a moral virtue. Second, the resemblance of the finely-grained beater trait with regard to more familiar coarsely-grained traits in child abusers also disposes us to think that the beater trait is not a moral virtue, because the coarsely-grained traits in child abusers systematically produce bad states of affairs in the actual world. This two-fold error theory explains away our intuitive reaction to the Mutor Case. Additionally, the error theory may be extended to undermine the intuition that benevolence is a morally better disposition than the beater trait. Perhaps the close correlation between "good" intentions and good consequences erroneously disposes people to think that "good" intentions enjoy moral value in themselves.

---

[9] Analogously, consider J. J. C. Smart's (1961) famous utilitarian substitute for moral blame. Because Smart does not attempt to provide an account of the pre-theoretical idea of blame, his view is widely rejected in the moral responsibility literature. See, for example, T. M. Scanlon (1988, p. 159), George Sher (2006, pp. 72-74), Angela Smith (2008, p. 374), and Manuel Vargas (2013, p. 166).

[10] Additionally, Michael Slote (2004, pp. 29-30) agrees that it is counterintuitive to grant that the beater trait is a moral virtue.

But is the error theory correct? Does an intuition infection prevent us from seeing that the beater trait is a moral virtue in the Mutor world and that benevolence and the beater trait are equally morally valuable in that world? Error theories are hard to refute, but they are also hard to corroborate. In view of that stalemate, a cost remains for an account of moral virtue that is unable to account for widespread pre-theoretical intuitions, because it is a cost for an account of a pre-theoretical idea if the account must explain away a core feature of it. Driver, however, has an additional response to supplement her error theory, which may eliminate that remaining cost.

Driver's second response aims to increase the plausibility of the claim that the beater trait is a moral virtue—and so that a disposition need not give rise to "good" intentions to be a moral virtue. She (2001, pp. 56-57) offers the following analogy:

> When parents take their child to the doctor to be vaccinated, they know that the shots will terrify the child and even cause the child pain. But they nevertheless are acting in the child's interest by getting her vaccinated as efficiently and competently as possible. The doctor's state of mind is irrelevant to them as long as it does not impact on the well-being of the child. The doctor who vaccinates the child may not have good intentions, in that the welfare of the child is not what is motivating him. He may only be motivated by his income, or the esteem of his colleagues, or the fact that the service is necessary for his certification. Yet, if he does a good job, he is the one parents will seek out.

Suppose that the doctor offers the most painless vaccinations in town and that the doctor's skill is the result of her selfish disposition. In that case, people in the community will value and choose that doctor over others. Analogously, those in the Mutor population also value the members of their community with the beater trait, which is a good ground for identifying the beater trait as a moral virtue.

The difficultly with the analogy is that there is no reason to think that the vaccination disposition is a moral virtue if the status of the beater trait as a moral virtue is itself being

questioned. That is, the very reason to deny that the beater trait is a moral virtue applies also to the doctor's vaccination disposition. In both cases, it is a disposition to form "bad" psychological states, which is sufficient to rule out its candidacy for being a moral virtue on a pre-theoretical basis. But perhaps I can say more. Driver appears to think that the social value of these traits alone is enough to grant them the status of moral virtue. The problem, however, is that a disposition's being socially valuable is not sufficient to indicate that it is a moral virtue. Driver (2001, p. 106) provides an example of such a case: "Suppose that someone smelled good, so good that whenever this person walked into a room, tensions were eased, conflict reduced, and so on." This perfume trait is socially valuable precisely because it systematically produces good states of affairs. It cannot, however, be a moral virtue, because it is a non-psychological trait. Nevertheless, the perfume trait may be a virtue in some other sense (Driver 2001, p. 106). It is instructive to note that Driver makes an implicit distinction between traits that are *mere virtues* and traits that are *moral virtues*. Both kinds of virtue systematically produce good states of affairs, but only psychological dispositions can be moral virtues.

I want to suggest that there is a more pre-theoretically plausible distinction in the same neighborhood. Perhaps the distinguishing feature between mere and moral virtues is not that only psychological dispositions are moral virtues but rather that only dispositions to form "good" intentions are moral virtues. On this revised distinction, the category of mere virtue includes non-psychological dispositions *and* "bad" psychological dispositions that systematically produce good in the world. This revised distinction not only maps onto our pre-theoretical intuitions about moral virtue, but it also partially accommodates Driver's claim that the beater trait and the vaccination disposition are moral virtues by allowing these dispositions to be virtues of some

inferior kind.[11] The upshot is that the doctor analogy fails to provide additional support for Driver's error theory, and a refined distinction between mere and moral virtues plausibly accounts for the socially valuable aspect of the beater trait and the vaccination disposition without admitting that they are moral virtues.

Aristotelian and Kantian accounts of moral virtue are able to accommodate the relevant intuitions and principles. On both of their views, moral virtues are necessarily dispositions of the mind and will worthy of admiration. For Aristotle (2002, 1105a30-5; 1120a23), a moral virtue is a disposition to aim at what is noble or fine for its own sake, and, for Kant (1996b, 6:405), a moral virtue is a disposition to do one's duty out of respect for the moral law. Thus, both accounts rule out dispositions to form "bad" intentions from being moral virtues. Furthermore, both accounts imply that dispositions with psychology worthy of admiration possess superior moral value to equally instrumentally valuable dispositions with "bad" psychology. For Aristotle (2002, 1107a5-1107a15), dispositions with "bad" psychology are bad in themselves and thus are morally inferior to dispositions with psychology worthy of admiration. For Kant (1996a, 4:394), psychology worthy of admiration is constitutive of intrinsic moral value, and thus a disposition to form an intention to act according to the moral law out of respect for it is morally superior to a disposition with "bad" psychology. Thus, one sympathetic to Driver's account cannot dismiss these intuitions and principles as unproblematic for her account because rival accounts of moral virtue cannot endorse them either.

[11] My revised distinction between mere and moral virtues appears to be the same as Michael Stocker's distinction between *technical* and *moral* virtues (Driver 2001, pp. 57-58). Technical virtues systematically produce good states of affairs but lack the "good" intentional and motivational structure required to be a moral virtue. They are virtues of an inferior sort. Driver's (2001, p. 58) main response to Stocker is that his distinction either commits a person to virtue maximalism where only the best traits can be moral virtues or it is a distinction without a function. I agree that virtue maximalism is to be avoided. So then, given the parallel between my and Stocker's distinctions, is it also incumbent upon me to show that the revised distinction between mere and moral virtues serves some function? My answer is that my distinction serves the same function as Driver's implicit distinction between mere and moral virtues. Both distinctions aim to accommodate pre-theoretical intuitions.

Arguments from Luck

Driver's account is committed to *actualism*—the view that the supervenience base for the moral status of a disposition is limited to the consequences that it produces in a single possible world. This commitment makes her account vulnerable to several objections from luck. The first objection highlights the counterintuitive way in which luck may affect which dispositions are counted as moral virtues or vices. Ben Bradley (2005, pp. 292-294) offers three versions of this objection, which we may call instances of *the Fragility Argument*.

First, suppose there is a violent psychopath named Lucky who is the only malicious person in a particular world, and his malice never produces harm. As it happens, Lucky coughs as he begins to insult others, birds fly in the path of his deftly aimed bullets, and his cellphone alerts potential victims of his presence. Since the moral status of malice is exhausted by the consequences it produces in that world and malice produces no bad consequences, malice is not a moral vice in that world. Thus, fluky events can counterintuitively affect the moral valence of a disposition. Second, suppose that there is some happy world where no one is in need and thus there is never an opportunity to exercise generosity. In such a world, generosity produces no good states of affairs, and so generosity is not a moral virtue in that world (cf. Driver 2001, p. 81). Thus, lack of opportunity to exercise a disposition may counterintuitively affect its moral valence. Third, suppose that super-benevolence is a uninstantiated property in some world. Since uninstantiated properties do not cause anything, super-benevolence produces no good states of affairs and thus is not a moral virtue in that world.[12] It is counterintuitive, however, that

[12] There is an even more radical version of the Fragility Argument that we may call *the Inversion Argument* (cf. Driver 2001, pp. 79-80; Slote 2004, p. 30; Calder 2007, pp. 204-208). Instances of the Fragility Argument show several ways in which luck may prohibit traits that are pre-theoretically moral virtues or vices from being such, and the Inversion Argument takes it a step further illustrating how luck may render a psychological trait that is pre-theoretically a moral virtue to be a moral vice.

dispositions which happen not to be instantiated in some world have no moral valence. The case intuitions in Bradley's Fragility Argument are captured by this general principle: necessarily, no disposition is a moral virtue or vice due to luck. And on Driver's account, there are several different ways in which luck affects the moral status of a disposition.[13]

Let us now consider a new objection from luck against Driver's account. This objection aims to show that whether an agent is a good or bad person at a particular time may counterintuitively depend on geographic luck in some substantial sense. Call this objection *the Location Argument*. Even though a world is a limit on the consequences that count toward the moral status of the disposition, there may be more than one *context* or non-overlapping supervenience base within a world (Driver 2001, p. 85).[14] A disposition may systematically produce good states of affairs in one actual context and bad states of affairs in another, and thus a psychological disposition may be a moral virtue in one actual context but not in another. Moral virtues that easily take on another moral status in a different context are called 'fragile' moral virtues (Driver 2001, p. 37). For example, Jane Bennett's blind charity is a moral virtue in rural England, because it makes her pleasant to accompany (Driver 2001, p. xvii). Her blind charity, however, would be a moral vice if she were to occupy Nazi Germany, because it disposes her not to notice the injustices perpetrated by the Nazis.

Suppose that Jane Bennett is not only blindly charitable but is also modest, trusting, and forgiving. Driver defines these virtues of ignorance as follows:

---

[13] Someone sympathetic to Driver's account may attempt to counterbalance the negative impact of the Fragility Argument by appealing to the attractiveness of the actualist requirement itself. In a later section, however, I will attempt to demotivate a commitment to actualism.

[14] The lines dividing supervenience bases are metaphysically vague. This, of course, results in metaphysically vague moral virtues and vices.

- Blind Charity: Jane focuses on the good qualities of others and is ignorant of their defects (2001, p. 28).
- Modesty: Jane is ignorant but not badly ignorant of her own worth (2001, p. 19).
- Trust: Jane believes p in spite of great evidence against p (2001, p. 31).
- Forgiveness: Jane forgives and forgets the wrongs done to her and others (2001, p. 32).[15]

In the context of rural England, all four of these dispositions are fragile moral virtues. Suppose that Jane Bennett travels from rural England to Germany during Nazi supremacy. Due to that shift in context, the supervenience base of the moral status of her traits changes. This provides an opportunity for the moral status of her traits to change. Not only does blind charity become a moral vice, but modesty, trust, and forgiveness also become moral vices. Being modest, Jane would not value her own opinion enough to challenge Nazi authority. Being trusting, Jane would believe that the German government is working for the common good against her best evidence. Being forgiving, Jane would forgive and forget the wrongs that the Nazis have done to her and others. The result is that Jane may be morally virtuous in her old context but morally vicious by unwittingly stepping into a new context. We may even suppose that Jane's trip to Germany is only the first stop in a series of seven visitations. If each stop is a new context, the moral status of Jane's character traits may change with each stop, and Jane may have a different *moral* character every day of the week.

Furthermore, it is not the case that Jane's character traits must be enacted in her new context for their moral valences to shift, because Driver's account is a trait-*type* and not a trait-*token* systematic account.[16] The difference between those systematic accounts concerns the scope of consequences included in the supervenience base of the trait. On the trait-type view, the

[15] Owen Flanagan (1990) and Michael Winter (2012) argue that these traits are not moral virtues at all.

[16] The distinction between trait-token and trait-type systematic production is the same as Driver's (2001, p. 79) atomistic and systematic distinction.

consequences that *all instances* of generosity produce within a particular context count toward the moral status of each person's generosity, but, on the trait-token view, only the consequences that a *particular token* of generosity produces count toward its moral status. As a result, the moral status of Jane's dispositions depends on the consequences that other tokens of her trait-types produce in this new context. This implies that Jane's particular traits need not be enacted for them to gain a new moral status.

Whether an agent is a good or bad person is determined by the moral valences of her dispositions. A person who is overall morally virtuous is a good person, and a person who is overall morally vicious is a bad person. But then, if the moral status of an agent's dispositions changes as she moves into different contexts, then the quality of her moral worth (whether she is a good or a bad person) may also change. In rural England, Jane may be a good person, but, in Nazi Germany, she may be a bad person. The counterintuitive result is that someone may become a better or worse person by inadvertently moving into a new contextual location. This kind of constitutive moral luck seems to be even more counterintuitive than the sort of constitutive moral luck that Thomas Nagel (1979, pp. 32-33) identifies: an agent may be praiseworthy or blameworthy for possessing an involuntarily acquired disposition.[17] The following general principle explains the intuition that Jane's moral worth should not change as she moves into new spatial locations: necessarily, an agent's being a good or bad person does not depend in any substantial sense upon where she is at a particular moment.

[17] One might think that Nagel's (1979, p. 28) category of circumstantial luck is a better fit than the category of constitutive luck for the phenomenon that I identify. The problem is that circumstantial luck pertains to the morally significant *decisions* one faces. More specifically, the issue is that an agent might choose better or worse depending on the proper subset of possible circumstances in which she finds herself. In contrast, constitutive luck is about the praise or blame one deserves for the moral status of her dispositions. Thus, constitutive luck is the right category for this phenomenon.

How might one sympathetic to Driver's account respond to the Location Argument? One might concede that being a good, neutral, or bad person is context-dependent but not in the counterintuitive way that I suggest. That is, one might reject the account of moral worth wherein the moral quality of a person is exhausted by the actual moral status of her character traits. Instead, one might affirm a characterization of moral worth that is determined by the actual moral status of her character traits as well as the moral status that her character traits would have if she were to occupy each actual context. On this view, Jane's moral worth is determined by the conjunction of her actual and counterfactual moral worth for all actual contexts. Call this alternative view of moral worth *the subjunctive account of moral worth*. An implication of this view is that Jane's moral worth does not change as she moves in and out of various contexts even if the actual moral status of her traits do change, because the actual moral status of her dispositions does not exhaust her moral worth. The subjunctive account of moral worth thus avoids the counterintuitive implication of the Location Argument.

But is the subjunctive account of moral worth plausible? There are at least two reasons not to think so. The first is that the proposal itself is counterintuitive. Asking the question 'what kind of person is Jane?' intuitively admits of answers such as 'a good person', 'a bad person', or 'neither a good or bad person'. In contrast, the subjunctive view of moral worth answers the question in this way: Jane is a good person relative to rural England, a bad person relative to Nazi Germany, and so on for all actual contexts. Relatedly, answering the question 'what kind of person is Jane?' intuitively appeals only to the particular traits Jane possesses and how they *actually* supervene but not additionally how they *counterfactually* supervene in contexts Jane is not occupying. In other words, it is counterintuitive that Jane's moral worth may be shaped by moral virtues and vices that she possesses only in a counterfactual sense.

The second reason is that Jane's moral worth may still be significantly affected by luck. By dividing the actual world into different supervenience bases, each supervenience base may be far more restrictive than the world limit, which may allow luck to have a very significant impact on the moral status of dispositions. As a result, there may be some actual contexts in which there are not many instances of the same type of traits Jane possesses, and, in those contexts, her traits may take on a moral valence quite different from the moral valence of her traits in her usual contexts. In that case, the content of Jane's subjunctive moral worth may be substantively enhanced or tainted by fluky events, which seems like the wrong result. The upshot is that the subjunctive account of moral worth avoids the counterintuitive implication of the Location Argument only by being counterintuitive in different ways.

Because other major accounts of moral virtue are not subject to either objection from luck, these are not difficulties that Driver's account can escape on account of the fact that they are problems for everyone. Aristotle's (2002, 1097a-1098b) account avoids the Fragility Argument, because whether a disposition is a moral virtue or vice is grounded in human nature. As a result, even if a disposition luckily causes uncharacteristic effects, fails to be actualized, or fails to be instantiated, the disposition retains its status as a moral virtue, moral vice, or neither.[18] Aristotle's account also avoids the Location Argument, because it is plausibly not a contextual account. And since Kant's (1996b, 6:405) account of moral virtue includes no consequentialist element, it evades both the Fragility and Location Argument.

[18] Driver (2001, p. xv) interprets Aristotle's theory of moral virtue as a mixed theory: "certain psychological states are necessary for virtue but not sufficient, since a virtue trait must show some connection to actual human flourishing." It is worth noting that Todd Calder (2007, pp. 213-219) plausibly argues that Aristotle's account of moral virtue is an intrinsic and not instrumental view. The exercise of a moral virtue is itself constitutive of (and does not instrumentally cause) happiness or flourishing.

The Modified Account

Driver's account is less plausible than it could be. The problem lies with the commitment to *multiple context actualism*. On that view, only those consequences that a disposition produces within a possible world count toward its moral status—and a world may include distinct supervenience bases. In contrast, *counterfactualism* is the view that a trait's supervenience base includes the consequences it produces within a world and in other close possible worlds (Driver 2001, p. 78). That is, a moral virtue is a psychological disposition that systematically produces good states of affairs in a world and close worlds. I will argue that Driver's account of moral virtue becomes far more pre-theoretically plausible if it is recast as a counterfactual account. Call this reformulated account of moral virtue *the modified account*. The modified account maintains the basic integrity of Driver's account of moral virtue, because the only change that it makes is to widen the scope of consequences that count toward the moral status of the disposition. In what follows, I will show the way in which the modified account mitigates or eliminates each difficulty from the previous two sections.

First, the modified account more nearly captures the pre-theoretical idea of moral virtue than Driver's original account with respect to this principle: necessarily, no moral virtue is a disposition to form "bad" intentions. The reason is that the modified account is able to rule out more dispositions to form "bad" intentions from being moral virtues than Driver's account. To see this clearly, it is helpful to distinguish between two types of disposition that systematically produce good. The first type of disposition systematically produces good states of affairs in a *lucky way*, and the second type of disposition systematically produces good states of affairs in a *non-lucky way*. Appealing to Duncan Pritchard's (2014) modal account of luck, a lucky event is

an event that occurs in some world but not in a wide class of possible worlds close to it.[19] As a result, dispositions that luckily systematically produce good states of affairs do so in the actual world but not in a broad range of close worlds, and dispositions that non-luckily systematically produce good states of affairs do so in the actual world and in a wide class of close worlds.

The modified account rules out the possibility of dispositions to form "bad" intentions that luckily produce good states of affairs from being moral virtues. Suppose that malice luckily and regularly produces good in the actual world. In that case, it is plausible that malice systematically produces bad states of affairs in most close possible worlds. But then, given the modified account, malice is not a moral virtue in the actual world where it regularly produces good states of affairs, because it regularly produces bad states of affairs in most possible worlds relevant for determining its moral status. This result generalizes. Since the modified account rules out one type of disposition to form "bad" intentions from being moral virtues that Driver's account allows to be moral virtues, the modified account more nearly captures the pre-theoretical idea of moral virtue.

Of course, even though the modified account comes closer to the pre-theoretical idea of moral virtue, the modified account cannot endorse the principle that necessarily, no moral virtue is a disposition to form "bad" intentions. For on the modified account, all dispositions to form "bad" intentions that systematically produce good states of affairs *in a non-lucky way* are still moral virtues—just as they are on Driver's account. With regard to this class of "bad" disposition, the modified account and Driver's account are on par, because the non-lucky

[19] In earlier work, Prichard (2005, p. 128) asserts that luck not only has a modal condition but also has a significance condition. Prichard (2014, pp. 604-606) now believes that including the significance condition is a mistake. It is worth noting that Prichard's modal account of luck may easily be extended to account for the intuitive idea that luck comes in degrees (cf. Church 2013, pp. 39-42).

mechanism that fixes a disposition's systematic production of good states of affairs in a world also fixes it that way in a broad range of close worlds.

Second, the modified account is able better to accommodate the claim that for two equally instrumentally valuable dispositions in a particular world, a disposition with psychology worthy of admiration is morally superior to a disposition to form "bad" intentions. Recall that there are two relevant categories of "bad" disposition—those that produce good in a lucky or non-lucky way. The modified account offers an explanation for why an admirable disposition is morally superior to a "bad" disposition that luckily produces good states of affairs: a disposition worthy of admiration produces more good in close possible worlds than does a "bad" disposition that only luckily produces good. The reason is that a disposition worthy of admiration produces good by the non-lucky mechanism of "good" intentions, and thus this disposition systematically produces good in most close possible worlds. Nevertheless, for the same reason as before, the modified account and Driver's account are on par with regard to admirable and "bad" dispositions that each produce good in a non-lucky way.

Third, the broader supervenience base in the modified account prohibits dispositions from taking on their moral status in a lucky way. Even if Lucky's malice fails to produce harm by some fluke, he nevertheless succeeds in bringing about the harm that he maliciously intends in worlds close to that one. As a result, malice is a moral vice in the world in which it luckily produces few (or no) bad results. Furthermore, supposing that there is no opportunity to exercise generosity in a particular world, there are plausibly opportunities to be generous in close possible worlds. Because generosity is likely to produce good consequences in those worlds, it is plausible that generosity is a moral virtue in the world in which it is instantiated but never actualized. Of course, if generosity produces no good consequences in close possible worlds,

then it is not the case that generosity fails to be a moral virtue by luck. Moreover, it is plausible that at least some traits that are not instantiated in a possible world are instantiated and activated in close worlds. As a result, at least some traits that are not instantiated in a particular world may yet be moral virtues or vices. The modified account is thus able to protect dispositions from luckily gaining their moral status in any of these ways.

Fourth, the modified account is not subject to the Location Argument. On the modified account, there is only one supervenience base of a disposition's moral status, and its scope contains the consequences a disposition produces in a world and close worlds. The upshot is that the moral worth of an agent cannot change as she changes locations. Furthermore, the modified account of moral worth avoids the problems with the subjunctive account of moral worth. The agent's moral worth is determined exclusively with regard to the way in which her psychological dispositions *actually* supervene—and not also how her dispositions *would* supervene had things been different.[20] Furthermore, the modified account of moral worth is not subject to the influence of fluky events on the agent's moral worth, because the wider supervenience base plausibly rules out any substantial influence of luck. Thus, the modified account has several advantages over Driver's original account.

Demotivating Actualism

With all of these advantages in adopting counterfactualism, why choose actualism in the first place? What theoretical fruitfulness lies in actualism that makes it more appealing than counterfactualism? Driver (2001, p. 81) offers the following argument for actualism:

[20] To be clear, moral virtues and vices on the modified account *actually* supervene on consequences that dispositions produce in the actual and close possible worlds. It is not the case that moral virtues or vices *counterfactually* supervene.

> My primary problem with the counterfactual approach is that it favors nearby possible worlds or privileges them. . . . Suppose, for example, that Sally would have had bad traits if she had not been raised by her mother, who, it turns out, did raise her only through amazing luck – the mother was almost run over by a truck but avoided death through an amazing fluke. Well, in worlds very close to this one, Sally is a bad person. Her high spirits become disruptive. Does it make sense to say that her high spirits are a vice in our world because in worlds close to ours, though not in ours, they are disruptive?

Sally's spiritedness systematically produces good states of affairs in the actual world, because the guidance and discipline of her mother habituates her spiritedness in a particular way. In close possible worlds where her mother dies, however, Sally has not received the benefit of moral training, and her spiritedness systematically produces bad states of affairs. Driver's intuition is that Sally's spiritedness in the actual world is a moral virtue, which is supposed to be an intuition that supports actualism over counterfactualism.

On the contrary, this thought experiment provides no support for actualism. In order for the thought experiment to favor actualism over counterfactualism, Sally's spiritedness must pick out the same disposition in the actual world as it does in close possible worlds, and the moral status of the disposition must be more intuitive given actualism than given counterfactualism. The problem is that the meager details of the thought experiment plausibly indicate that Sally's spiritedness is not the same disposition in the actual world as it is in close possible worlds. The habituating influence of Sally's mother produces a different variant of spiritedness in the actual world than Sally has in close possible worlds where her spiritedness is not so habituated. But now we have a satisfying explanation for the different moral valence of her spiritedness in different worlds: Sally's spiritedness is a different disposition in these worlds. In other words, the

problem with Driver's thought experiment is that it is actually a case of constitutive luck—luck regarding which variant of spiritedness Sally possesses.[21]

Nevertheless, Driver (2001, p. 81) also offers a related thought experiment that does have the potential to support actualism over counterfactualism: "Suppose that high spirits produce good in the actual world through some cosmic fluke, yet fail to in nearby possible worlds. Sally's high spirits are still virtuous." Additionally, Driver (2001, p. 82) offers this brief argument to support the intuition that Sally's spiritedness is a moral virtue:

> The argument for … [actualism] is simply to point out that if virtue evaluation is to serve a meaningful function in providing evidence of a person's reliability in producing good, … [actualism] makes a lot of sense. What happens in nearby possible worlds has little practical significance to this world, and moral evaluation does serve a practical function.

In Driver's view, if a disposition *x* systematically produces good states of affairs in the actual world due to luck (and so not also in close possible worlds), *x* is plausibly a moral virtue in the actual world, because attributing the status of moral virtue to *x* serves the practical function of highlighting a person's reliability in producing good.

Interestingly, however, Driver's reasoning about the practical purpose of virtue evaluation appears to support counterfactualism over actualism. The purpose of moral virtue evaluation is to flag dispositions that reliably produce good. Actualism confers the status of being a moral virtue to dispositions that either *reliably* or *luckily* systematically produce good states of affairs. But a disposition that systematically produces good states of affairs via a cosmic

[21] My analysis of the problem with Driver's argument differs from Bradley's diagnosis (2005, p. 291). According to him, the difficulty is that Sally's spiritedness does not itself produce good consequences in the actual world and bad consequences in nearby possible worlds. Rather, he argues that her spiritedness produces good or bad consequences in concert with other dispositions which differ between the actual world and close possible worlds. Bradley's analysis of Driver's argument, however, ends here.

fluke in the actual world is not a reliable good-producing disposition. Thus, actualism may bestow the status of moral virtue to a disposition that does not reliably produce good. In contrast, counterfactualism grants the status of moral virtue *only* to dispositions that reliably systematically produce good states of affairs, because its broader supervenience base eliminates the possibility that a disposition gains its moral status by luck. As a result, there is good reason to think that counterfactualism better serves the practical purpose of virtue evaluation, which is to flag reliable good-producing dispositions.

One might think that this response flounders because it invokes a different sense of reliability than the one to which Driver plausibly appeals. On the *frequency* sense of reliability that Driver invokes, a reliable good-producing disposition is identical with a disposition that actually systematically produces good—whether the good is produced in a lucky or non-lucky way. That is, a reliable good-producing disposition is merely a disposition with an actual track record of regularly producing good. On the *propensity* sense of reliability that the objector employs, a reliable good-producing disposition is identical with a disposition that has an intrinsic structure which makes it very likely that the disposition regularly produces good in the actual world. On this view, it is possible but very unlikely that a reliable good-producing disposition fails to produce good systematically in the actual world.

Even though the objector employs a different sense of reliability, there is still a substantive question about which sense of reliability is the important one for moral evaluation. I will argue that the propensity sense of reliability better serves the practical purpose of moral evaluation than the frequency sense in at least one way. Suppose that people around Sally know that the intrinsic structure of her spiritedness is such that its actual track record of producing good is a mere fluke. Looking to the future, no one would plausibly trust that Sally's spiritedness

will continue to produce good, because there is no good reason to think that this streak of luck will continue. (Of course, this is not the case if the disposition turns out to produce good via a non-lucky mechanism.) Since one practical purpose of moral evaluation to flag dispositions that reliably produce good in the future, the frequency sense of reliability *only sometimes* serves this practical purpose. It allows unreliable good-producing dispositions to be moral virtues. In contrast, the propensity sense of reliability that counterfactualism employs is a far better candidate to serve this function meaningfully, because the propensity sense marks good-producing dispositions as reliable only if they are likely to produce good in the future. Thus, counterfactualism has the advantage in practical reasoning that there are always good objective reasons to rely on moral virtues to produce good, whereas there are not always good objective reasons to rely on moral virtues if actualism is true.[22]

In addition to counterfactualism's advantage in practical reasoning over actualism, counterfactualism also has a significant metaphysical advantage if the future does *not* exist. Compare two dispositions that produce an equal amount of past and present good in the actual world. One disposition is an intrinsically good-producing disposition, and the other disposition is a lucky good-producing disposition. If the future exists, actualism grants a better moral status to the intrinsically good-producing disposition than the lucky good-producing disposition, because the intrinsically good-producing disposition brings about more future good. If, however, the

[22] An analogy illustrates this point nicely. Suppose there are two racehorses, Thunder and Lightning. In Thunder's practice runs, his times are consistently worse than other racehorses. Nevertheless, Thunder has won all five races in which he has competed. At the end of each race, all of the leading horses trip and fall, which allows Thunder to take the lead and win. In contrast, Lightning's practice times consistently beat the times of the other horses. It is clear that Lightning is stronger and faster than the other horses. Lightning too has won all five races in which he has competed. Since these horses have the same winning record, their good-producing track record is equivalent. According to the frequency sense of reliability, these horses are equally good. Nevertheless, looking to the future, which horse would you rely on to win the race? Obviously, the rational bet is Lightning. And it is still the rational bet even if Lightning had won only four out of five of her previous races. The analogy offers intuitive evidence that it is the propensity sense of reliability that is the important sense for the practical purpose of flagging reliable good-producing dispositions, which is a contention that supports counterfactualism.

future does not exist, then actualism evaluates the intrinsically good-producing disposition and the lucky good-producing disposition in the same way, which is counterintuitive. In contrast, if the future does not exist, counterfactualism grants the intrinsically good-producing disposition a better moral status than the lucky good-producing disposition. The reason is that the intrinsically good-producing disposition brings about far more good in close possible worlds. So then, if the future does not exist, counterfactualism has a significant metaphysical advantage over actualism, and, whether or not the future exists, counterfactualism has an advantage in practical reasoning over actualism.

Moral Virtue and Right Action

There is a further issue about how well the modified account of moral virtue fits with accounts of right action. After all, if the modified account of moral virtue does not plausibly complement any utilitarian account of right action and the original account does, this provides a good reason to retain the original account.[23] Consider the view of right action that Driver (2012, pp. 112-130) endorses: an action is the right action if it produces the most good and/or least bad consequences in the actual world given the options available to the agent.[24] Her accounts of right action and moral virtue are theoretically symmetric, because, in both cases, the relevant moral property is determined exclusively by consequences produced in the actual world. Theoretical symmetry is a good-making feature of a systematic ethical theory, because the motivations behind each element reinforce the motivations behind the other elements. For example, the general reasons to adopt Driver's account of moral virtue typically also apply to her account of right action.

[23] I thank an anonymous referee for offering this objection.

[24] I ignore a complication concerning the distinction in practical reasoning between actualism and possibilism regarding which options are relevant for determining what action is best to pursue. For an overview of this subject, see Driver (2012, pp. 131-144).

Does the modified account of moral virtue also complement Driver's account of right action? The answer is no. Some of the reasons that favor the modified account of moral virtue have analogues that count against Driver's account of right action. For example, the Fragility Argument draws out a counterintuitive feature of Driver's account of moral virtue, and its analogue has the same impact on her account of right action. The analogue would highlight the counterintuitive affect that luck may have on the morality of an action: a malevolent act may turn out to be the right action to perform if it luckily produces the most good given the options available to the agent. So then, the modified account of moral virtue and Driver's account of right action do not fit well together into a systematic ethical theory, because they can stand together only in a theoretical house divided against itself.

There are, however, consequentialist accounts of right action that suitably complement the modified account of moral virtue. Consider the *foreseeability account* of right action: "the right action is the action that produces the most foreseeable good and/or the least bad [in the actual world] amongst options open to the agent at the time of action" (Driver 2012, p. 112). As Driver (2012, p. 112-113) notes, foreseeable is not identical with foreseen. The agent's mental states do not constitutively determine which action of a range of options is the right action. A plausible way to unpack the idea of foreseeability is to appeal to what an ideal observer would foresee. An ideal observer has access to all the relevant evidence in order to correctly calculate the objective probability that each option has for producing good. The right action typically has a high objective probability of producing overall good results.[25] Interestingly, this significantly mitigates the possibility of an act gaining its moral status in a lucky way, because an act with the high intrinsic probability of producing good is sure to produce good in a wide class of close

[25] I hedge with 'typically' because there may be cases where the option with the highest objective probability of producing good in fact has a low objective probability of producing good or preventing bad.

worlds. This modal implication of the foreseeability account of right action allows it to fit nicely with the modified account of moral virtue. Moral virtues always produce good in a wide class of close worlds, and right actions almost always produce good in a broad range of close worlds.

The account of right action that is theoretically symmetrical with the modified account of moral virtue is the *counterfactual account* of right action: an action is the right action if it produces the most good and/or least bad consequences in the actual world and close worlds given the range of options available to the agent. The counterfactual account of right action and the modified account of moral virtue have similar supervenience bases. In both cases, the moral property is determined by the consequences produced in the actual world and close worlds. Adopting this account of right action has many of the same benefits of endorsing the modified account of moral virtue. For example, the counterfactual account of right action protects the moral status of the agent's action from being affected by luck. If an action produces the most good in a wide class of possible worlds but luckily fails to produce good in the actual world, it is still the right action, because it is the best good-producing option. Thus, the counterfactual account of right action plausibly integrates with the modified account of moral virtue into a systematic ethical theory.

Of course, there is no space here to evaluate the plausibility of these three accounts of right action, and they do not exhaust the relevant options. But if either the foreseeability or counterfactual accounts of right action is plausible, then it is not the case that the modified account cannot fit with any plausible utilitarian account of right action. Furthermore, the attractiveness of either the foreseeability or counterfactual accounts of right action may add to the attractiveness of the modified account of moral virtue—and vice versa. If, however, Driver's account of right action is more plausible than the foreseeability or counterfactual accounts, this

consideration counts in favor of Driver's original account of moral virtue. Thus, even if Driver's original account of moral virtue is less plausible than the modified account, there may be systematic ethical reasons that privilege adopting the original account over the modified account.

Conclusion

Driver's account of moral virtue is subject to various objections from admiration and luck. In particular, I argued that Driver's account cannot endorse these principles:

(1) Necessarily, no moral virtue is a disposition to form "bad" intentions.
(2) Necessarily, a disposition with psychology worthy of admiration is morally superior to a disposition with "bad" psychology for two equally instrumentally valuable dispositions in the same world.
(3) Necessarily, no disposition is a moral virtue or vice due to luck.
(4) Necessarily, an agent's being a good or bad person does not depend in any substantial sense upon where she is at a particular moment.

The modified account better captures the pre-theoretical idea of moral virtue with regard to each of these principles. In particular, the modified account comes closer to the pre-theoretical idea of moral virtue behind principles (1) and (2) than Driver's account—even though neither account can endorse those principles. Regarding (1), the modified account precludes more dispositions to form "bad" intentions from being moral virtues than Driver's account. Regarding (2), the modified account implies that the class of admirable good-producing dispositions is morally superior to the class of "bad" dispositions that luckily produce good. Driver's account does not have this implication. Furthermore, the modified account endorses principles (3) and (4), and Driver's account is incompatible of them. And since one measure of plausibility for an account of moral virtue concerns how well it captures the pre-theoretical idea of moral virtue, there is good reason to think that the modified account is more plausible than Driver's original account.

Additionally, there is at least one way in which the modified account better serves the practical purpose of moral virtue evaluation than Driver's original account. Moreover, if the future does not exist, the modified account has a metaphysical advantage over Driver's account. Therefore, anyone sympathetic to Driver's account of moral virtue has several good reasons to adopt the modified account.

Even so, depending on which account of right action turns out to be correct, there may be overriding reasons for a consequentialist to retain Driver's original account of moral virtue and absorb the costs. In other words, even if the modified account is more plausible than the original account of moral virtue, there may be advantages in retaining Driver's original account in order to preserve theoretical purity. Conversely, if the counterfactual account of right action is the most plausible utilitarian account, theoretical purity is an additional reason to favor of the modified account of moral virtue.[26]

[26] I am grateful to Joel Archer, Josh Anderson, Donald Bungum, Gideon Jeffrey, Anne Jeffrey, Max Perish, Josh Rasmussen, Nick Setliff, and several anonymous referees for their comments on some version of this essay. I owe a special debt of gratitude to Elizabeth Foreman for her comments on many versions of this paper. Additionally, I am thankful for the commentary provided by Julia Driver at the Illinois Philosophical Association, Brad Cokelet at the Eastern Division of the American Philosophical Association, and Scott Forschler at the Indiana Philosophical Association. I also thank the audience members for their questions and comments.

REFERENCES

Aristotle. (2002). *Nicomachean ethics*. Christopher Rowe (Trans.), Oxford: Oxford University Press.

Bradley, B. (2005). Virtue consequentialism. *Utilitas, 17*(3), 282-298.

Calder, T. (2007). Against consequentialist theories of virtue and vice. *Utilitas, 19*(2), 201-219.

Church, I. M. (2013). Getting 'lucky' with Gettier. *European Journal of Philosophy, 21*(1), 37-49.

Driver, J. (2001). *Uneasy virtue*. Cambridge: Cambridge University Press.

Driver, J. (2004). Response to my critics. *Utilitas, 16*(1), 33-41.

Driver, J. (2012) *Consequentialism*. New York: Routledge.

Flanagan, O. (1990). Virtues and ignorance. *The Journal of Philosophy, 87*(8), 420-428.

Kant, I. (1996a). *Groundwork of the metaphysics of morals*. In M. J. Gregor (Ed. and Trans.), *Practical philosophy* (pp. 39-108). Cambridge: Cambridge University Press.

Kant, I. (1996b). *The metaphysics of morals*. In M. J. Gregor (Ed. and Trans.), *Practical philosophy* (pp. 353-604). Cambridge: Cambridge University Press.

Nagel, T. (1979). Moral luck. In T. Nagel (Ed.), *Mortal questions* (pp. 24-38). Cambridge: Cambridge University Press.

O'Neill, O. (2004). Consequences for non-consequentialists. *Utilitas, 16*(1), 1-11.

Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.

Pritchard, D. (2014). The modal account of luck. *Metaphilosophy, 45*(4-5), 594-619.

Ross, W. D. (1930). *The good and the right*. Oxford: Clarendon Press.

Scanlon, T. M. (1988). The significance of choice. In S. M. McMurrin (Ed.) *The tanner lectures on human values* (Vol. 8, pp. 149-216). Cambridge: Cambridge University Press.

Sher, G. (2006). *In praise of blame.* Oxford: Oxford University Press.

Slote, M. (2004). Driver's virtues. *Utilitas, 16*(1), 22-32.

Smart, J. J. C. (1961). Free will, praise, and blame. *Mind, 70*(279), 291-306.

Smith, A. (2008). Control, responsibility, and moral assessment. *Philosophical Studies, 138*(3), 367-392.

Vargas, M. (2013). *Building better beings: A theory of moral responsibility.* Oxford: Oxford University Press.

Winter, M. (2012). Does moral virtue require knowledge? A response to Julia Driver. *Ethical Theory and Moral Practice, 15*(4), 533-546.