



Is Fair Opportunity a Comprehensive Theory of Responsibility?

Douglas Husak¹

Accepted: 2 May 2023

© The Author(s), under exclusive licence to Springer Nature B.V. 2023

Abstract

I challenge the adequacy of David Brink’s “master principle” of culpability. I allege that it fails to account for the moral relevance of ignorance of wrongdoing. I describe three cases in which I believe that Brink’s theory of normative competence cannot account for the significance of a variable that bears on culpability. In most of this paper I attempt to anticipate and reply to the various responses Brink might offer to my challenge.

Keywords Responsibility · Ignorance · Culpability · Blameworthiness

1 Introduction

David Brink has done an extraordinary job identifying, defending, and applying what might be called a “master principle” of culpability. He has written a wonderful book, full of nuance and philosophical sophistication, and nothing I say here is intended to contest the quality of his impressive effort. In the best philosophical tradition, however, I will express my central reservation in what follows.

According to his (deceptively) simple master principle, “responsibility is a matter of an agent having the requisite capacities, which we can conceptualize in terms of *normative competence*, and suitable opportunities to exercise the capacities free from external interference of certain kinds” (FO&R, p.4). An agent A is blameworthy for his wrongful conduct Φ iff A had a fair opportunity to avoid Φ -ing. In order to have such an opportunity, A must have adequate degrees of cognitive and volitional competence. I assume Brink holds that agents A and B must be equal in blameworthiness for their same act Φ iff they possess the same cognitive and volitional capacities with respect to Φ -ing as well as the same opportunities to exercise them. That is, given comparable opportunities, any factor that marks a contrast between their degrees of responsibility for Φ -ing must be a matter of their

✉ Douglas Husak
husak@philosophy.rutgers.edu

¹ Rutgers University, New Brunswick, USA

different amounts of cognitive or volitional competences. Although earlier and far less detailed versions of this general view were famously defended by HLA Hart, no one has provided a more complete and careful explication of it.

My main reservation with Brink's position is that his master principle fails to account for the significance of a factor I hold to be crucial to determinations of blameworthiness, especially once we allow (as does Brink) that responsibility is scalar and admits of degrees. I will begin by describing what I take to be a class of counterexamples to his master principle—a class of cases in which I believe that agents A and B are *not* equal in their degree of blameworthiness when they commit the same wrongful act Φ , even though they possess identical degrees of cognitive and volitional competence as well as comparable opportunities to exercise them. In most of my paper, I will examine possible ways that Brink might respond to my objection—even when these responses deviate from what he explicitly says. I do not believe that any of these responses is likely to succeed. I conclude that we must qualify or supplement his master principle if we hope to produce an adequate theory of responsibility. But I admit to uncertainty. Ultimately, we should not be too confident that we understand exactly what is subsumed within an agent's volitional and cognitive capacities. The line that separates what a wrongdoer *could not have* done from what he *could* but *didn't* do, has bedeviled moral and legal theorists since philosophy began.

Elsewhere I have described the kind of examples I believe are problematic for Brink's master principle. They involve pairs of cases in which two agents, A and B, commit the same wrongful act Φ and differ only in their beliefs about the moral status of Φ . In short, A and B commit the same act Φ , and A knows Φ is morally wrong but B does not; B believes Φ is permissible. In each of the following pairs of cases, I stipulate (if it is coherent to do so) that A and B are equal in their amount of cognitive and volitional competence as well as in their opportunities to exercise their capacities. I claim that in nearly every such case (the qualification is unimportant for present purposes) A is *more* blameworthy than B. Awareness of wrongdoing, I claim, makes A more responsible for Φ than B. If, however, A turns out to be more blameworthy than B in these cases, and Brink's master principle is correct, there must be something amiss with my stipulation; there must be some difference in the normative or cognitive competence of A as opposed to B. I do not believe any such difference in the cognitive or volitional competence of A and B can be found, but it is hard to be certain. Perhaps a difference can be detected if a more thorough analysis of cognitive or volitional competence can be provided.

I need examples of the kind of cases I have in mind. Generating such cases should be easy; situations in which one agent does what he knows to be wrong while another agent acts in moral ignorance are plentiful. Because I caution that not too much emphasis should be focused on any particular example, I offer three separate pairs of cases that satisfy my general description. First, the pair of cases that has garnered the most philosophical commentary involves slave-owning at different historical periods. Suppose A and B both engage in the same act(s) Φ , the act of owning slaves. B, an ancient Hittite, is unaware that slave-owning is wrong. The wrongness of the institution of slavery has not even occurred to B. A, another slave-owner, knows full well that slave-owning is wrong but cannot be bothered to do the hard

work his slaves perform on his plantation. A is from the same historical era as B, but has somehow managed to grasp the moral horror of slavery. I hold that A is more blameworthy for Φ than B. Perhaps B is not blameworthy *at all* for owning slaves, but I need not defend this radical position. For present purposes, I assert only that B's moral ignorance renders him *less* blameworthy than A for Φ .

In order to avoid some of the distraction that seems to have surrounded this first example, I offer an additional pair of cases I believe to share the same structure (despite a few differences to be discussed below). Suppose A and B are adult males who desire a spouse to bear their children. To do so, both participate in an institution of arranged marriage in which they pay the father of their prospective bride a sufficient amount of money to induce him to give his consent to the subsequent sexual relations. The bride herself is not consulted about the marriage and actively opposes it. She emphatically withholds her consent to mate with her arranged husband. A knows this arrangement is morally wrong but engages in it anyway because he does not want to take the trouble to persuade a woman to marry him. By contrast, B believes arranged marriages are morally permissible; consent to sex is not required from women as long as their fathers approve. I believe that B's moral ignorance about the institution of arranged marriages renders him less blameworthy than A for Φ .

Third and finally, suppose A and B are diners who patronize an exotic foreign restaurant and order from a menu. Both select octopus for their hors d'oeuvre and monkey for dinner. A knows these animals are highly intelligent and that eating them is morally wrong. Still, he enjoys their taste and allows himself this immoral indulgence. B believes it is permissible to eat any non-human. I contend that A is more responsible than B for Φ . If I am correct that A is more blameworthy than B for owning slaves, for participating in the sex that accompanies an arranged marriage, or for eating intelligent animals, knowing these acts are morally wrong, it follows (given my stipulation) that there are variables that contribute to relative degree of blameworthiness other than their cognitive and volitional competence. Thus Brink's master principle is mistaken; responsibility for wrongdoing is not simply a function of normative and cognitive capacities and opportunities to exercise them.

What does Brink actually say about moral ignorance and its potential to create an excuse? His comments are made in two different places. In an instructive footnote, he cites my work as well as that of several other philosophers, and admits that "it is reasonably common to think that moral ignorance can excuse" (58, n14). He continues: "While I agree that non-culpable ignorance excuses, I disagree with those who think that awareness is a requirement of culpability" (58, n14). I gather he also believes that awareness is immaterial to the *relative* culpability of A and B. Thus Brink contends that the sole difference between A and B—their differing awareness that what they are doing is wrong—is immaterial to their degree of blameworthiness. On my view, by contrast, moral ignorance itself precludes or at least mitigates blameworthiness for Φ . Thus our disagreement is stark. In the remainder of this paper, I will discuss several replies Brink might make (some of which he *does* make) to my challenge. Of course, I believe the best way to respond is to revise or supplement his master principle and allow that factors other than fair opportunity,

analyzed in terms of cognitive and volitional competence, are variables that contribute to judgments of blameworthiness.

First, Brink might challenge my judgments about the moral status of Φ itself. I concur with Brink that questions about the blameworthiness of agents for Φ -ing ordinarily arise only when Φ -ing is morally wrong. My first two pairs of cases are deliberately chosen to make this response implausible. To be sure, practices of slavery, arranged marriages, and the eating of intelligent animals have been common throughout most of human history. Nonetheless, it seems extraordinary to deny that these acts are wrong; acts of owning slaves or having sex with women without their consent are as clear examples of wrongs as can be given. Any philosopher who dares to suggest that slavery or non-consensual sex might be permissible risks immediate cancellation from his progressive colleagues. Perhaps relativists of a certain stripe might try to contend that these acts can be permissible in some circumstances, but (I am happy to report) nothing in Brink's book leads me to believe he is tempted by relativism or would challenge my moral judgments about Φ in my first two pairs of cases (although I suppose he could deny that slavery as practiced in one time and place is a "different act" from slavery as practiced elsewhere). To settle the matter, he explicitly writes "slavery is a profound wrong" (296). Of course, the third pair of cases is different; the basis for believing Φ is wrong in this example is more contentious, as there is far less consensus today about the wrongfulness of eating even the most intelligent non-human animals. I select this third example as a case in which future generations may look back, as we do with institutions of slavery and arranged marriage, and wonder how agents in an earlier time could have been so misguided.

Second, Brink might simply deny the truth of the judgments of responsibility I make in the above pairs of cases. Perhaps A and B *are* equally responsible; their different beliefs about the moral status of Φ are immaterial to their quantum of blameworthiness. I believe that this is the response Brink actually makes to my challenge—at least in the first of my three pairs of cases. Ultimately, I have only a little to say against philosophers who do not share my intuitive judgments about these cases; conflicts of moral intuitions are notoriously hard to resolve. I support my own judgments by relying on the very Strawsonian reactive attitudes Brink regards as a pretty reliable guide to whether and to what extent agents are blameworthy. Respondents may answer a survey by saying that moral ignorance does not excuse or mitigate. But I think that when most respondents are told about how A and B differ in their beliefs about slavery, arranged marriages or eating intelligent animals, their judgments of the relative degree of blameworthiness of A and B differ accordingly. In case there is doubt, imagine that *you* start off as agent B in the above hypotheticals. Despite your initial state of moral ignorance, you are said to be as responsible as A for owning slaves, having sex with a woman with whom you have mated despite her lack of consent, or eating intelligent animals. I am confident you would respond to this accusation of blameworthiness by protesting that you were unaware that what you were doing is wrong. You may (or may not) hold yourself to have been somewhat blameworthy for Φ before you transitioned to agent A and became convinced of its moral wrongfulness, but I strongly doubt you would hold yourself to be *as* blameworthy as A, who continues to Φ even after he learns the inconvenient moral truth.

I bolster my conclusion with a real anecdote. Let us suppose that the rule against using a phone in an area explicitly prohibited by Homeland Security is designed to minimize risks of coordination between terrorists or other wrongdoers at national borders. A few years ago, I observed a stranger using a mobile phone in such an area. It is easy to anticipate how the offender reacted when an authority confronted him. He certainly did not reply, “I know what I am doing is wrong, but I hoped not to be caught.” Nor did he reply, “I have nothing to say on my own behalf; ignorance of wrongdoing is no excuse.” Instead, he responded, “I am sorry; I did not know I was not allowed to use my phone here.” I make two observations if I am correct to assume that this latter reply is nearly universal and the former two are highly unusual. First, the offender must have believed he was entitled to leniency—some degree of exculpation—if his plea were accepted as true. Second and just as importantly, his plea of ignorance is often *accepted* as wholly or partially exculpatory by the authority who confronted the offender. One would be surprised to learn that this offender did not actually *receive* some degree of leniency, and probably no sanction whatever, relative to that of an offender who knew the use of mobile phones is prohibited but hoped to avoid detection for his wrong. If the intuition that ignorance is no excuse were as entrenched as some philosophers allege, we should be puzzled by the fact that ordinary persons plead it so frequently and authorities accept it so readily (unless they believe offenders are lying). But these facts are *not* puzzling. A comprehensive perspective on the blameworthiness of ignorant wrongdoers must explain rather than neglect these truisms. Our Strawsonian reactive attitudes, I submit, create problems for Brink’s master principle.

Notwithstanding this support for my intuitions, Brink himself disagrees with my judgments about the relative blameworthiness of A and B. He spends a fair amount of time discussing my first example of slave-owning (although not my latter two examples). If his master principle is to generate an excuse for the Hittites, these ancient wrongdoers must have lacked the amount of normative competence possessed by those who are aware slave-owning is wrong. Indeed, Brink is somewhat tempted by this position, even though he ultimately rejects it. There is no reason to believe that ancient Hittites lacked normative competence across the board. Incompetence, however, “can be selective” (294). The Hittites might have been suffering from what Brink calls a “moral blindspot” (294ff). He writes: “Those who are not generally incompetent can have *moral blindspots* that lead them to engage in wrongdoing that is selective but nonetheless robust in character. I’m thinking especially of people who treat some people but not others unjustly out of a principled moral mistake about the moral standing and rights of others with whom they interact” (295). The interesting cases for Brink are those in which the agent’s impermissible behavior “is the product of moral blindspots that the agent learned and internalized from his family and community that are deeply entrenched” (295). Ancient Hittites are good candidates to suffer from these blindspots; they would not dream of enslaving other Hittites but do not respect the rights of those they have conquered in battle to live as free men and women.

Even if they suffered from a blindspot, Brink is unsparing in blaming the Hittites. He admits that they grew up in societies in which “the practice of slavery was widespread and, we might suppose, widely believed as permissible” (296). Still,

he does not believe that their moral ignorance about the practice of slavery is non-culpable, and he therefore rejects their plea for an excuse. As we have seen, Brink allows an excuse when moral ignorance is non-culpable, but not otherwise. When is moral ignorance non-culpable? His general answer, unsurprisingly, is that “ignorance is non-culpable only when the agent suffered a relevant cognitive incapacity” (58, n13). And the Hittites suffered no such incapacity. After all, “abolitionist ideas have been around as long as slavery” (297). Even Aristotle, who famously defended slavery, acknowledged the possibility that abolitionism was the preferable moral position. Brink does not require that given Hittites must “have been aware of moral challenges to slavery for their ignorance to be culpable” (297). As a “capacitarian,” he emphatically rejects that “awareness of wrongdoing is a condition of blameworthiness” (297 n47). What is needed to render the moral ignorance of the Hittites culpable is not their actual awareness of anti-slavery rhetoric, but their “capacity to recognize that slavery is wrong” (297). And the Hittites had this capacity. He writes: “If the wrongness of slavery can be grasped based on moral and non-moral beliefs available to the slaveholder exercising modest epistemic diligence, then the slaveholder is cognitively competent... All he needs to do is to see that the scope limitations on which rights he recognizes are arbitrary” (298).

I suspect that the use of the “blindspot” metaphor may be misleading. When we ask why otherwise competent individuals might be excused for their wrongful acts, it is tempting to look for a localized deficiency or shortcoming from which they suffer. In attempts to understand the difference in blameworthiness between A and B, however, it might be more instructive to focus not only on B, the morally ignorant wrongdoer, but also on A, the agent who engages in wrongdoing despite knowing he is doing so. What is it about *knowledge* that Φ is wrongful that could make A more blameworthy than B? Instead of trying to understand why *ignorance* of wrongdoing might *mitigate* blameworthiness relative to a baseline of *knowledge*, we might try to ascertain why *knowledge* of wrongdoing might *aggravate* blameworthiness relative to a baseline of *ignorance*. After all, moral ignorance about institutions of slavery and arranged marriage has probably been the norm rather than the exception throughout most of human history. The more interesting psychological question raised by my pairs of cases is not how people (e.g., B) could be so blinkered as to not comprehend the wrongness of what they were doing, but how people (e.g., A) could be so evil as to perpetuate institutions they knew full well to be horrific.

In any event, Brink continues to allow that Hittites could have a “blindspot,” even though he is unwilling to allow this condition to exculpate because “blindspots are corrigible” (299). His verdict on cases such as that of the Hittites confirms his “more general claim that when moral blindspots are selective they tend to be culpable because the agent already has underlying moral commitments in light of which the blindspot is corrigible” (298). He continues: “In effect, I’m arguing that in these particular cases of manifest immorality, moral blindspots don’t support a finding of genuine incompetence and, hence, don’t justify an excuse of selective incompetence” (299). I withhold judgment about whether Brink is correct that the moral ignorance of the Hittites about slavery is “corrigible,” by which I gather he means that they should have been able to “see through” and correct their lack of moral vision by exercising their existing cognitive capacities. I fear we may be asking too much of

these Hittites when we demand that they cultivate modern enlightened sensibilities and somehow figure out that slavery is wrong. But no matter. Actually, I concede that Brink is probably correct about the normative competence of the unknowing wrongdoers in each of my examples. In all likelihood, ancient Hittites had all of the competence they need to be responsible. I disagree, however, that this concession entails their lack of blameworthiness. I hold that agents can be excused from blame *despite* the fact that their normative and volitional competence is intact. Thus these examples provide excellent tests of whether Brink's master principle is correct.

But it is hard to be sure. The foremost difficulty in applying Brink's master principle to my pairs of cases is to provide a principled basis for differentiating between the following alternatives. Suppose A and B are identical in their capacities at time t_1 . At t_2 , however, A comes to learn something that B does not. Should we describe this transition by saying that A has gained a new capacity not possessed by B? Or should we describe this transition by saying that A has exercised in a different way a capacity he and B continue to share? Clearly, an answer to this question is crucial if we hope to establish how someone who holds Brink's master principle would begin to pass judgment on my hypothetical pairs of cases. If B lacks a capacity to understand that what he is doing is wrong, that is, if his blindspot is incorrigible, his ignorance is non-culpable and he cannot be blamed. But if A and B differ in how they exercise the capacities they both possess, a different quantum of blameworthiness can be attached to them without jeopardizing Brink's master principle. Unfortunately, no *general* answer to this question is likely to be available; some cases will fall on one side of the line, while some cases will fall on the other. Unless we have some way to identify which alternative is correct, however, I am pessimistic that Brink's master principle will allow us to make judgments about particular cases in which our intuitions are unclear. A set of principles to draw this elusive line might be said to constitute the *gold standard* in applying his master principle to pairs of cases.

Without the gold standard, easy cases will remain easy, whereas hard cases will remain hard. Suppose (to use Brink's example) that A learns to speak Russian. It seems fair to claim he has gained a new capacity relative to B: A can speak Russian; B cannot. But suppose A looks at a thermometer and learns the truth of a proposition of which B is unaware, viz., the temperature is 70 degrees. It would be strained (to say the least) to claim that A has gained a new capacity relative to B. Instead, A has merely gained new information by exercising a capacity he and B share in common. I speculate that these examples are clear because we can say that when A learns to speak Russian, he knows *how* to do something that B does not know how to do. When A reads the thermometer, by contrast, he knows *that* something is true that B does not. This example is easy because we take for granted that B *could* have done what A did to learn what A knows, that is, look at the thermometer to discover the temperature. To be able to do so, of course, B must know *how* to read a thermometer. If he lacks this ability, he suffers from the very deficiency in capacities that Brink allows to be exculpatory.

Epistemologists since at least the time of Gilbert Ryle have debated whether and to what extent it is important to distinguish knowing *how* from knowing *that*. Is the former a distinctive kind of non-propositional mental state? Or is knowing *how*, like

knowing *that*, a species of propositional knowledge? I hope not to take a firm side on this well-worn but complex topic. My only point is that this contrast might serve as a rough guide to whether A has gained a new capacity relative to B, as opposed to whether A has merely exercised a capacity he and B share. Here, then, is the beginning of a principled basis for the gold standard needed to help resolve pairs of cases involving moral ignorance: When A learns *how* to do something he previously did not know how to do, he gains a new capacity. But when he merely learns *that* a proposition is true about which he was previously ignorant, A does not gain a new capacity but simply exercises a pre-existing capacity.

The foregoing suggestion is helpful, but I doubt that it will take us very far. The contrast between knowing *that* and knowing *how* is too easily gerrymandered to resolve disputes about blameworthiness. Still, if this suggestion has *any* merit in developing the gold standard we seek, the question before us is how to apply it to the pairs of cases I have constructed. If we imagine that A and B begin with the same moral knowledge, but A comes to know that Φ is wrong while B reaches no such insight, should we say that they differ in their exercise of their common capacities, or should we say instead that A has gained a new capacity? I take Brink to suppose that in my pairs of cases both A and B have the same capacities as well as similar opportunities to exercise them. Here I tend to agree. But Brink thinks it follows from this supposition that A and B must be equal in their degree of responsibility for Φ -ing. Here I tend to disagree. A and B *do* differ in their degree of responsibility; B's ignorance is at least a partial excuse. Thus Brink's master principle gives us (what I intuit to be) the incorrect judgment in the kind of cases I have been examining.

To his credit, however, Brink is characteristically modest and cautious, and anticipates a line of inquiry he might have pursued further. He sensibly entertains the possibility that a given blindspot could be incorrigible because it results from the fact that "someone [is] brought up in a very insular and homogeneous community and tradition who is taught to privilege the needs and interests of members of her community in relation to the needs and interests of outsiders" (299). Patty Hearst is "arguably" an example of someone who has been indoctrinated. The lore about Patty, almost certainly exaggerated, is that she was "brainwashed" (more precisely described by Brink at p.267 as a state of "temporary non-culpable cognitive incompetence" induced by another). As a result, she failed to appreciate the wrongfulness of her complicity in murder and robbery. I suppose it is correct that "brainwashing" (whatever it may be) is a good candidate for a condition that would impair capacities and provide an excuse.

Again, however, I hope not to become too immersed in any particular example. Let me apply some of the foregoing modest and cautious remarks to the two remaining pairs of cases I have constructed. My question is how Brink *could* (even if he *would* not) analyze the second and third of my examples. My second case, it should be recalled, involves a man who does not believe he needs the consent of a woman to have sex because he has bought her from her father. Do those males who participate in institutions of arranged marriages possess the competence to grasp the immorality of the practice by which they obtain supposed consent to sexual relations? To be sure, no one would hold males to be "brainwashed" when they believe

that they need not garner the consent of their spouse to mate, as long as the father of the bride has consented. But Brink's reservations do not require cases of incorrigible blindspots to involve anything like the brainwashing alleged to have taken place with Patty Hearst. They only require, in his words, that a male participant in an arranged marriage be "brought up in a very insular and homogeneous community and tradition [and] is taught to privilege the needs and interests of members of her community [that is, males] in relation to the needs and interests of outsiders [that is, females]." This description of the circumstances of male participants in arranged marriages seems apt. Perhaps, then, they suffer from an incorrigible blindspot they cannot be blamed for failing to see though. But does this analysis also apply to my final pair of cases in which a diner eats very intelligent animals? Can we describe this diner as experiencing a blindspot when he does not realize he is acting wrongly by eating octopus and monkey? Has he been indoctrinated? I think not. Perhaps my reluctance to apply these labels, however, results from my own proximity to a culture in which eating intelligent animals is a pervasive practice. We have a very hard time recognizing our own indoctrination.

I assume there must have been some moment of time in my third example when A came to believe what B does not, namely, that eating intelligent animals such as octopus and monkey is wrong. When A made this transition about Φ , should we say he acquired a new capacity—as if he had learned *how* to speak a new language? I believe it is more plausible to say he simply employed the capacity he already possessed and *that* he reached a new insight about the permissibility of eating animals. Humans do not develop new capacities whenever they change their moral beliefs and judge an act to be wrong which they previously held to be permissible. Brink is probably correct that no new capacities are developed by those who come to realize the impermissibility of slavery or non-consensual sex via an arranged marriage. Or so it seems to both Brink and me—although I admit to being less than certain. I confess that I have no easy way to furnish the details of a gold standard that would allow us to differentiate between cases in which A has gained a new capacity not possessed by B, as opposed to those in which A exercises in a different way a capacity they both share. Without this gold standard, it will be difficult to apply Brink's master principle to help us decide whether A and B differ in their quanta of blameworthiness.

Where do we stand? When a slave-owner who never questions the morality of the institution changes his mind and adopts a whole new perspective on the universal rights of man, is it really so absurd to say he has gained a new capacity? When someone accustomed to the institution of arranged marriages comes to realize that the consent of a woman is more transformative than that of her father, is it really so outrageous to conclude that an entirely different capacity has been created? When someone comes to believe that non-human intelligent animals have a right not to be eaten, is it really so crazy to think that a new capacity has come into existence? After all, these are not trivial changes in moral opinion; they are tectonic shifts in an agent's whole moral outlook.

I am not sure, however, whether there is a *fact of the matter* about how these questions should be answered. Perhaps *moral* considerations rather than metaphysical speculations, should dictate the outcome. In light of this uncertainty, I offer a

tentative proposal about how to proceed. In a very different context, Brink entertains the possibility that “fallible appraisers in the real world” (389) might be forced to resort to non-ideal theory to resolve hard problems. If philosophical reflection and institutional mechanisms cannot be trusted to draw elusive boundaries sensibly, we need a default principle in order to move ahead. Many candidates are available. I prefer a principle of *lenity* according to which reasonable doubts about blameworthiness should be resolved in favor of the accused: When in serious doubt, don’t blame. This principle is more familiar in legal contexts, but there is just as much reason to apply lenity to the moral realm. The state of being blameworthy is dreadful, and should not be conferred lightly. Brink could have retained his master principle and concurred with my moral intuitions by holding that new capacities are created when wrongdoers make profound changes in their fundamental moral outlook.

I invite Brink to embrace the line of thought I have just developed. More likely, however, he will stick to his original judgment about the Hittites and hold that moral ignorance is almost always corrigible and therefore rarely excuses. If so, A and B are equally blameworthy in each of my three pairs of cases. A third alternative, almost certainly even more unpalatable to Brink, is to supplement or revise his master principle and allow that moral ignorance can have independent force as at least a partial excusing condition.

Acknowledgements I would like to thank Kim Ferzan for helping me to avoid errors in a previous draft of this essay.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.