

CRIME, BLAMEWORTHINESS, and OUTCOMES

James Edwards* and Andrew Simester**

Abstract—If criminal law blamed in a way that accurately reflected blameworthiness, what would it say about the outcomes of our actions? On one view, criminal law would be outcome-*insensitive*: all crimes would be defined in the inchoate mode, and outcomes would be irrelevant to the quantum of punishment. On a second view, criminal law would be doubly outcome-*sensitive*: some crimes would be defined in terms of outcomes, and more punishment would be imposed where they occurred. Here, we reject both of these views in favour of a third. While the outcomes of our actions affect which wrongs we commit, they do not make us more blameworthy for committing them. A criminal law that accurately reflected blameworthiness would convict those who commit significantly different wrongs of different crimes. It would punish those who are equally blameworthy to the same degree. So the outcomes of our actions would be relevant to criminalisation. But they would be irrelevant to the quantum of punishment imposed.

Keywords – Blameworthiness; Crime; Moral Luck; Punishment; Wrongdoing.

1. Introduction

Criminal law's responses to crime, it is commonly claimed, are blaming responses. By convicting and punishing D, the criminal law blames D, or communicates that D is to blame, for criminal conduct.¹ In life outside the law, blaming responses are subject to the ethics of blame – to norms that bear upon instances of blame, and upon

* Associate Professor of Law, University of Oxford. Email: james.edwards@law.ox.ac.uk.

** Provost's Chair in Law, National University of Singapore and Edmund-Davies Professor of Criminal Law, King's College London. Email: simester@nus.edu.sg. We presented an early draft at the Criminal Conversations workshop, in London, at the end of June 2017. Thanks to the participants on that occasion; also to Antony Duff, Alex Sarch, Findlay Stark, Bob Sullivan and two anonymous OJLS reviewers for written comments.

¹ See e.g. J Gardner, 'Wrongs and Faults' in AP Simester (ed) *Appraising Strict Liability* (OUP 2005) 77; D Husak, *The Philosophy of Criminal Law* (OUP 2010) 347-350; AP Simester and A von Hirsch, *Crimes, Harms and Wrongs* (Hart 2011) 11-13; V Tadros, *Wrongs and Crimes* (OUP 2016) 160-162. Note too the title under which Michael Moore presents his theory of criminal law: M Moore, *Placing Blame* (OUP 1997).

its communication. There is no reason to think that criminal law's blaming responses are exempt from these norms. No doubt there are many additional considerations that bear on whether D should be convicted and punished. But the considerations supplied by the ethics of blame do not disappear when we turn to legal life.² Or so, at least, we will assume here.

Suppose that D acts wrongly. One aspect of the ethics of blame is concerned with potential blamers: with who may blame D, or communicate that D is to blame.³ A second aspect is concerned with D's blameworthiness. Consider the following norm, which for convenience we call (Blaming):

(Blaming): *Ceteris paribus*, blaming responses should accurately reflect blameworthiness. A blaming response to D's wrong accurately reflects blameworthiness if, and only if, D is blamed for the wrong in proportion to D's blameworthiness for committing it.

² Much debate in criminal law theory can be made sense of only on the assumption that this is so. To take just one example, consider criminal defences. Why is so much time spent asking, with reference to ever more complicated examples, whether D's conduct is morally justified or excused? Why is it plausible for Mitch Berman to claim that on the 'standard account' of justificatory and excusatory defences, these categories 'simply replicate or mirror the distinction that obtains in moral theory'? The obvious answer appeals to the ethics of blame: if D is not blameworthy, a blaming response is out of place. We are not blameworthy for conduct that is morally justified or excused. All else being equal, a defence should therefore be available. For Berman's discussion, see his 'Justification and Excuse, Law and Morality' (2003) *Duke Law Journal* 1, 7.

³ See e.g. GA Cohen, 'Casting the First Stone: Who Can, and Who Can't, Condemn the Terrorists?' in his *Finding Oneself in the Other* (Princeton University Press 2013); RA Duff, 'Blame, Moral Standing and the Legitimacy of the Criminal Trial' (2010) 23 *Ratio* 123.

To endorse (Blaming) is not to say that blameworthy conduct should always be met with a blaming response, let alone with criminal conviction or punishment. It is merely to say that *if* there is to be such a response—and all else is equal—that response should accurately reflect D's blameworthiness. Because of the *ceteris paribus* clause, (Blaming) is compatible with the conclusion that, once countervailing considerations are factored in, conviction and punishment typically should not reflect blameworthiness.⁴ The truth of this conclusion would not render (Blaming) unimportant. (Blaming) would continue to give the countervailing considerations something they must defeat.

Our aim here is not to defend (Blaming) or the common assumptions with which we began.⁵ It is rather to explore some of their implications. If they are correct, there are reasons for criminal conviction and punishment to accurately reflect the blameworthiness of the convicted and punished. Setting aside additional considerations, we can then ask what a criminal law that conformed to those reasons would look like. Our focus here is on one part of the picture. It is on what such a criminal law would have to say about the outcomes of what we do—about the results and consequences of our actions.

⁴ What we say here is therefore compatible with a variety of views about the justification of punishment. As David Enoch and Andrei Marmor put the point, punishment is one of many *blame-related reactions*. A complete justification of these reactions would make reference both to blameworthiness, and to various other considerations: 'The Case Against Moral Luck' (2007) 26 *Law and Philosophy* 405, 412-13. Our (and their) concern is with the former.

⁵ We say something about why people should be blamed for wrongs in section 5. One explanation of why we should blame those who act wrongly in proportion to their blameworthiness is that doing so is part of giving them what they deserve.

2. Outcomes in Criminal Law: Three Views

Let us start with a beguilingly simple answer: absolutely nothing. A criminal law that accurately reflected blameworthiness, some say, would be comprehensively insensitive to outcomes. Most notably, it would be insensitive both when we criminalise and when we punish. Accordingly:

(C-I): All crimes would be defined without reference to outcomes.

(P-I): The quantum of punishment imposed by criminal courts would be unaffected by outcomes.

Perhaps the most distinguished defender of this view is Andrew Ashworth.⁶ For Ashworth, it is what we try to do that makes our actions wrong, and that makes them the wrongs they are.⁷ Whether we succeed is, for these purposes, beside the point. Consider:

⁶ A Ashworth, 'Taking the Consequences' in S Shute, J Gardner and J Horder (eds) *Action and Value in Criminal Law* (OUP 1993). For a more recent defence, see L Alexander and KK Ferzan, *Crime and Culpability* (CUP 2009). Joel Feinberg endorses a similar view in 'Equal Punishment for Failed Attempts: Some Bad but Instructive Arguments Against It' (1995) 37 Arizona LR 117, albeit without developing a positive justification.

⁷ For ease of exposition, we focus—like Ashworth—upon the distinction between completed attempts and successes. Unintended acts of endangerment present analogous issues. We say something about them in section 7.

Bowmen: D1 and D2 are bowmen standing at opposite ends of a field. V happens to be walking through the field in between them. D1 and D2 loose arrows in the direction of V, each intending that their arrow kill V. Both arrows are on course for their target. D1's arrow is intercepted by a passing duck. D2's arrow punctures V's chest, killing V instantly.

Ashworth argues that D1 and D2 commit the very same wrong: that of trying to kill V. They should be blamed for the very same thing, and they are equally blameworthy for it. In a system of criminal law that accurately reflected blameworthiness, D1 and D2 would be convicted of the same crime and suffer the same amount of punishment. The same is true of all pairs of defendants who act wrongly and differ only in that one succeeds while the other fails. So we should accept both (C-I) and (P-I).

To defend these claims, Ashworth borrows a suggestion made by Thomas Nagel.⁸ It is a precondition of moral assessment, Nagel suggests, that we have sufficient control over the object of the assessment. The problem with outcomes is that there are always factors that bear on their occurrence which we do not control. Bowmen have no control over passing ducks. Because 'the only element of an action which lies within the agent's control is the trying',⁹ all outcomes are ineligible for moral assessment. The control condition, as it is often called, is not met. Ashworth

⁸ T Nagel, 'Moral Luck' in his *Mortal Questions* (CUP 1979).

⁹ Ashworth (n 6) 108.

concludes that a criminal law which accurately reflected blameworthiness would have nothing to say about the outcomes of our actions.

This, at least, is Ashworth's defence of (C-I) and (P-I). Two important objections have been pressed against it. Both can be put in the form of a *reductio*. The first *reductio* was hinted at by Nagel himself, and developed in detail by Michael Moore.¹⁰ Suppose we accept, with Ashworth, that the control condition precludes blameworthiness for outcomes because uncontrolled factors always bear on their occurrence. Such factors do not only bear on whether we succeed. They also bear on whether we try. One reason for this is what whether we try depends on the circumstances in which we find ourselves. Had V not happened to walk through the field, D1 and D2 would not have tried to kill her. The same is true if D1 and D2 had been struck by lightning the moment before V appeared. Or if D1 and D2 had grown up before the invention of bows and arrows. To generalise the thought, we lack control over many of our circumstances, many of which bear on whether we try. Obviously enough, uncontrolled factors also bear on which choices we make, and on the character traits that influence our choices. The control condition implies that we cannot be morally assessed for any of these things—for our successes, our endeavours, our choices or our character. Pursued to its logical conclusion, it implies that we cannot be morally assessed *for anything*. In Nagel's words, the area of 'legitimate moral judgment, seems to shrink under this scrutiny to an extensionless

¹⁰ Moore (n 1) ch 5.

point'.¹¹ But there are surely some cases in which we *are* eligible for moral assessment. So Moore concludes that the control condition must be rejected.¹² We can add that the condition does nothing to support (C-I) and (P-I). Far from implying that we should revise the criminal law to make it insensitive to outcomes, the control condition implies that *no* blaming response accurately reflects blameworthiness. All else being equal, (Blaming) then implies that no-one should be convicted and punished.

This first *reductio* targets Ashworth's defence of the claim that D1 and D2 are to blame for the same wrong. A second *reductio* targets the claim itself. This objection was hinted at by Tony Honoré, and developed in detail by John Gardner.¹³ On Ashworth's view, our moral agency extends to our endeavours, but not to what those endeavours cause to occur. This is why those who succeed and fail commit the same wrong. Notice however that the moral case for trying is, at least standardly, derivative. We have moral reason to try to save a life, because we have moral reason to save it. We have moral reason not to try to kill, because we have moral reason not to kill. We have no moral reason to try to count grains of sand, because we have no moral reason to count them. And so on. Now if our moral agency does not extend to

¹¹ Nagel (n 8) 35.

¹² The condition might be weakened. One might say that we have sufficient control over an action if we can reliably convert the choice to so act into that action. This will not help those who wish to claim that outcomes are ineligible for moral assessment. We can often reliably convert the choice to bring about an outcome into that outcome. So this weaker interpretation of the control condition renders us eligible for moral assessment in respect of some outcomes.

¹³ T Honoré, 'Responsibility and Luck' (1988) 104 LQR 530, 543; J Gardner, 'Obligations and Outcomes in the Law of Torts' in P Cane and J Gardner (eds) *Relating to Responsibility* (OUP 2001). Gardner helpfully juxtaposes the two *reductios* in his 'Wrongdoing by Results: Moore's Experiential Argument' (2012) 18 Legal Theory 459.

our successes, there are no moral reasons to succeed. We just saw that our moral reasons to try are, at least standardly, parasitic on such reasons. So the non-existence of moral reasons to succeed entails, at least in standard cases, the non-existence of moral reasons to try.¹⁴ But it is surely the case—indeed, it is crucial to Ashworth’s case—that there *are* such reasons: if there are no moral reasons not to try to kill, there can be no moral *wrong* of trying to kill. Once we accept—as we surely should—that there are moral reasons not to try to kill, we must also accept the moral reasons not to kill from which they derive.¹⁵ And if there are such reasons—if our moral agency does extend to our successes—our motivation for denying that there is a wrong *of killing* falls away.¹⁶ Ashworth’s claim that D1 and D2 merely commit wrongs of endeavour, Gardner concludes, is a claim we should reject.

The force of these objections has convinced many to reject both (C-I) and (P-I). On a rival view, things are dramatically different. If criminal law is to accurately reflect blameworthiness, it *must* be sensitive to outcomes. This is true at both the criminalisation and punishment stages. In a criminal law that accurately reflected blameworthiness:

¹⁴ True, there are non-standard cases. We may have moral reason to try, even when we are doomed to fail. Attempting to save a loved one whose death is inevitable may be an expression of devotion (see Gardner, ‘Obligations and Outcomes’ (n 13) 38). Futile acts of protest may express commitment to a valuable cause. Let us grant that there are reasons to try in such cases. These reasons *can* survive the non-existence of reasons to succeed. Notice how odd it would be, however, to claim that our moral agency extends only this far: that in cases where apparent reasons to try are parasitic on apparent reasons to succeed, there are actually no reasons to try *at all*.

¹⁵ It might be said in reply that, though moral reasons not to try to kill are indeed derivative, they are derivative of the badness of killing, not of moral reasons not to kill. (We are grateful to both Alex Sarch and Antony Duff for raising this possibility.) To our mind, however, this response cannot stand. If killing is bad, there is something that counts against killing. And if there is something that counts against killing, there is reason not to kill.

¹⁶ Remember that appeals to control have already been ruled out.

(C-S): Some crimes would be defined with reference to outcomes.

(P-S): The quantum of punishment imposed by criminal courts would sometimes be affected by outcomes.

Defenders of this view include Michael Moore and John Gardner. For them, D1 and D2 in *Bowmen* commit different wrongs. They should be blamed for different things, and they are blameworthy to different degrees.¹⁷ In a system of criminal law that accurately reflected blameworthiness, D1 and D2 would be convicted of different crimes, and suffer different amounts of punishment. So we should reject (C-I) and (P-I), and accept (C-S) and (P-S).

In this paper, we stand up for a third view. On that view, those who try and those who succeed commit different wrongs, but the fact of their success does not itself make them more blameworthy.¹⁸ Because they committed different wrongs, D1 and D2 should – all else being equal – be blamed for different things. *Ceteris paribus*, they should be convicted of different crimes. But as they are equally blameworthy

¹⁷ For Moore's views, see Moore (n 1) ch 5, and M Moore, *Causation and Responsibility* (OUP 2009) ch 5. For Gardner's, see Gardner, 'Wrongdoing by Results: Moore's Experiential Argument' (n 13) and J Gardner, *Offences and Defences* (OUP 2007) ch 11.

¹⁸ Judith Thomson and Michael Zimmerman also endorse these claims: they distinguish between (i) the wrongs for which one is blameworthy and (ii) how blameworthy one is for those wrongs: J Thomson, 'Morality and Bad Luck' (1989) 20 *Metaphilosophy* 203; MJ Zimmerman, 'Taking Luck Seriously' (2002) 49 *Journal of Philosophy* 553. Both suggest that outcomes matter to the former but not the latter. Thomson defends these conclusions by appealing to two senses of blame. Zimmerman's argument appeals to a precondition of control. Below, we offer a different argument, which relies neither on the bifurcation of blame nor on the control condition.

they should – again, all else being equal – be punished to the same degree. This is to side with Ashworth on the quantum of punishment, and with Gardner and Moore on the definition of crime. We should accept (P-I) rather than (P-S), and (C-S) rather than (C-I). Or so we shall argue.

3. Outcomes in Morality: Wrongdoing, Blameworthiness and Engagement

Let us begin by distinguishing between some types of moral assessment. Inter alia, we can morally assess:

- (i) Actions
- (ii) Agents in respect of their actions
- (iii) Agents

One way to assess actions is to ask whether they are wrong. This is not a question we ask about agents. It makes sense to ask whether driving or singing are wrong, even if it is pretty obvious that they are not. It makes no sense to ask whether this is true of Bill or Sally.

What makes ϕ ing wrong? There are different views. Obviously enough, ϕ ing cannot be wrong if there is *no* reason not to ϕ . But many acts we have reason not to perform are not wrong. What must be added? Some say that the reasons not to ϕ must be *decisive*: they must defeat any reasons we have to ϕ . Others say that it is

wrong to ϕ only if the reasons not to ϕ are *mandatory*: only if that is, they exclude at least one reason to ϕ from consideration. Some add that the reasons against ϕ ing must be *categorical*: reasons which are immune to changes in our projects and goals. There are other possibilities.¹⁹ On all these views, however, what makes an action wrong is some property of the guiding reasons that bear on that action. So much, at least, seems to be a matter of general agreement.

Our main interest here is in assessments of type (ii). One way to assess an agent in respect of an action is to ask whether the agent is blameworthy for its commission. Suppose that D acts wrongly. D is blameworthy if and only if there is a particular type of *connection* between D and the wrong.²⁰ This is not a connection that follows automatically. Wrongdoing can be excused. And excuses preclude blameworthiness. In what, then, does the necessary connection consist? Crudely put, it consists in the fact that the action reflects badly on the agent – that it is appropriate to view D in a worse light in virtue of her wrong. When this is so, and only then, D is blameworthy for what she did.

When thinking about blameworthiness we can usefully distinguish between *preconditions* and *grounds*. Eligibility for assessment as blameworthy is determined by the relevant preconditions.²¹ If those preconditions are not satisfied, D is not blameworthy for having ϕ ed. But a world of eligible people might never be

¹⁹ For a range of views, see Tadros (n 1) ch 3.

²⁰ Our concern here is limited to blameworthiness for wrongs. Whether it is possible to be blameworthy for conduct that is not wrong is a further question that we cannot pursue here.

²¹ Two familiar examples in the criminal law are the requirements of voluntariness and sanity.

blameworthy for what they do. The grounds of blameworthiness make it the case that eligible actors are in fact blameworthy. They are considerations that, when present, establish the required connection between D and D's wrong: it is in virtue of their presence that D's conduct reflects badly on D.

We return to preconditions later. Let us first say a bit more about grounds. What makes it wrong to ϕ , we claimed above, is some property of the guiding reasons that bear on ϕ ing. What makes D's ϕ ing blameworthy, in contrast, is D's *engagement* with those guiding reasons. By engagement, we mean to refer to the way in which the reasons that bear on D's action figure in D's practical reasoning. D is blameworthy for ϕ ing in virtue of the fact that D's engagement with those reasons exhibits some shortfall or deficiency, relative to what we can reasonably expect from D. Call this *the engagement view*.

To see what we have in mind, imagine D shoots and injures V. The guiding reason not to injure V makes D's action wrong. That reason might figure in D's practical reasoning in a number of ways. D might have acted directly against it, by acting in order to injure V. It might be a reason D takes to count against shooting, but which D takes to be overridden by decisive reasons to shoot.²² Whatever the way, that (guiding) reason will figure in D's practical reasoning *for some further*

²² It might even be the case that D failed to pay the reason not to injure any attention at all. As noted earlier, our primary concern here is not with cases of unintended let alone inadvertent wrongdoing. We do not deny, however, that inadvertence can be blameworthy. For present purposes, we assume that one type of deficient engagement with reasons involves failing to attend to those reasons where one should. That negligence is blameworthy, then, is not ruled out by our analysis. Neither, for that matter, is it ruled in.

(*explanatory*) reason. Perhaps D acted in order to injure V because V thought this would be fun. Perhaps D took the reason not to injure to be outweighed because V's injury weakens a rival football team. These explanations do not, of course, reflect well on D. But there are others that paint D in a more flattering light. Perhaps D acted in order to injure V because she had been threatened with injury herself. Perhaps D took the reason not to injure to be outweighed, because on the facts available to D anyone would have concluded that injuring V was the only way to save many lives. Perhaps D paid no heed to the reason not to injure because D is unusually trusting, and T told D that no-one else was in the vicinity. These facts, and other like them, are standardly taken to determine whether D is blameworthy both within and without the law.²³ They are determinative, we suggest, because they tell us whether there was a deficiency in D's practical reasoning, one that we could reasonably expect D to avoid.²⁴ It is deficiencies of this kind, on the engagement view, that make D blameworthy – that constitute what we have called the grounds of blameworthiness.

One may ask what all of this has to do with outcomes. To see the answer let us return to *Bowmen*:

²³ They include facts that determine whether D's wrong was committed intentionally, recklessly or negligently, and whether D is entitled to a justificatory defence like necessity, or an excusatory defence like duress.

²⁴ Much is packed into this proviso that cannot be explored here. We might reasonably expect D not to act selfishly, but not expect D to avoid making mistakes that are understandable for one of D's intelligence. For discussion, see AP Simester, 'A Disintegrated Theory of Culpability' in D Baker and J Horder (eds), *The Sanctity of Life and the Criminal Law: The Legacy of Glanville Williams* (CUP, 2013) 178.

Bowmen: D1 and D2 are bowmen standing at opposite ends of a field. V happens to be walking through the field in between them. D1 and D2 loose arrows in the direction of V, each intending that their arrow kill V. Both arrows are on course for their target. D1's arrow is intercepted by a passing duck. D2's arrow punctures V's chest, killing V instantly.

What does the account we have sketched tell us about the blameworthiness of D1 and D2? It will help to separate three points in time.²⁵ At t_1 , D1 and D2 loose their arrows in the direction of V. At t_2 , D1's arrow is intercepted. At t_3 , D2's arrow punctures V's chest. That D1 and D2 both act wrongly is not contentious. Let us assume that both are eligible for blame, and have no justifications or excuses. What makes D1 and D2 blameworthy, we have suggested, is their engagement with the guiding reasons that bear on what they do. That engagement occurs at t_1 : D1 and D2 act in order to kill—both loose their arrows for reasons that, *ex hypothesi*, are defeated reasons. This exercise in practical reasoning is unaffected by what happens at t_2 and t_3 . Whether an arrow is intercepted, or ends up in V's chest, does not—indeed, cannot—alter the fact that D1 and D2 both acted for reasons that pointed decisively against what they did. Once the arrows are loose, it is too late for that. Because D1 and D2's engagement with the applicable guiding reasons is equally deficient, the engagement view holds that D1 and D2 are equally blameworthy for the wrongs they commit.

²⁵ We need to do this in the criminal law too. See the brief discussion of concurrence in section 7.

It is important to see that this remains the case if there are reasons to succeed. If there are such reasons, D1 and D2 commit different wrongs. D2 kills V. D1 only tries to do so.²⁶ This does not mean that there is a difference in how D1 and D2 engage with the applicable guiding reasons. *Ex hypothesi*, both have reasons not to kill. Both have reasons not to try to do so. The former reasons count against actions that increase the risk of killing, which helps explain why the latter reasons can standardly be derived from them. The key point here is that none of this drives a wedge between D1 and D2 when it comes to the grounds of blameworthiness. Both D1 and D2 engage with their reasons not to succeed, at t_1 , in the very same way: by trying to do what they have reason not to do successfully. Whether D1 and D2 actually succeed does not—indeed cannot—alter this exercise in practical reasoning: it cannot alter the way in which the applicable guiding reasons figured in their practical reasoning. By t_2 and t_3 , it is too late for that. So the conclusion reached in the previous paragraph remains: D1 and D2 are equally blameworthy for what they did.

These comments are neutral on the question of whether reasons (not) to succeed in fact exist. As it happens, we think they do. Our moral agency extends to our successes, as well as to our endeavours. The Honoré/Gardner *reductio* is one powerful reason to accept this. Another is found in the secondary moral duties we incur when we commit wrongs. The content of those secondary duties often turns on whether we succeed or fail. To see this, consider the following example. Dan and Ed

²⁶ These are different wrongs because they involve failing to conform to different types of guiding reason. Only D2 fails to conform to a reason not to kill, because only D2 kills V.

both throw stones in the street. They do so independently of one another. Dan's stone smashes Virgil's window. Ed's stone narrowly misses. Only Dan has a moral duty to compensate Virgil. Perhaps each of them has a duty to apologise, but the duty of compensation is Dan's alone: it exists *in virtue of* the fact that Dan caused Virgil to suffer harm. Ed, having caused no harm, has no such duty. Now it is hard to see why morality should impose extra secondary duties on Dan, if the fact that Dan actually harmed Virgil falls outside the scope of Dan's moral agency. That this duty exists, it seems to us, is further evidence that our moral agency is in fact more extensive. It is precisely because Dan committed a different wrong that Dan has different secondary duties to Ed.²⁷

What does not follow, however, is that we should see Dan in a worse light.²⁸ That Ed incurs fewer secondary duties as a result of throwing his stone in the street does not mean that we should think better of Ed for having thrown it. The same goes for our bowmen, D1 and D2. Like Dan and Ed, D1 and D2 committed different wrongs. By killing V, D2 committed a wrong that D1 did not. Doesn't it follow that D2 is more blameworthy? It does not. As we already saw, there is more than one

²⁷ Susan Wolf argues that we would expect Dan to offer to pay for the damage, and that we would be disturbed if he did not. She puts this down to a 'nameless virtue', which is 'a species of, or at least akin to, generosity'. While we agree that Dan can be expected to pay, we doubt that this has much to do with a hitherto unnoticed virtue, or with anything akin to generosity. Dan can be expected to pay, it seems to us, because he committed a different wrong, and because he therefore has different secondary duties. See Wolf, 'The Moral of Moral Luck' (2001) 31(1) *Philosophic Exchange* 1. Cf J Andre, 'Nagel, Williams and Moral Luck' (1983) *Analysis* 202, 205-206.

²⁸ Or that Dan should see himself in a worse light. Enoch and Marmor seem to assume that this does follow. Those who defend the moral significance of outcomes, they suggest, must think that outcomes alter our character for the worse, or that they make us more blameworthy: Enoch and Marmor (n 4) 409-411. Our suggestion here is that outcomes are sometimes morally significant for a different reason: because they make it the case that we commit additional wrongs. For the reasons given in the text that follows, this does not itself entail that we are more blameworthy, or that we are worse people.

type of moral assessment in play here. Conclusions about wrongdoing are the upshot of assessments of type (i). Conclusions about blameworthiness are delivered by assessments of type (ii). We cannot simply assume that the two types of assessment wax and wane together. Consider assessments of type (iii): assessments not of agents in respect of particular actions, but of agents *simpliciter*. Such assessments are often made in terms of virtue and vice. Some of us are kind and courageous and considerate. Others are cruel and cowardly and indifferent. Outcomes do not seem to bear on moral assessments of *this* kind. Imagine that T throws herself in front of V in an attempt to block D2's arrow. The courage T shows does not seem to depend on whom the arrow hits: if a gust of wind lifts D2's arrow over T's despairing dive, this does not make T any less courageous. Similarly, that D1's arrow is intercepted by a passing duck does not make D1 any less indifferent to V's life than D2. That we are courageous or indifferent people is a function of the courage or indifference we exhibit on occasions like these. If this is right, there are at least some assessments of type (iii) that are outcome-insensitive. The point here is not that it follows that assessments of type (ii) are similarly insensitive. The point is rather that, if assessments of type (i) are outcome-sensitive—as we have claimed they are—it does *not* follow that assessments of type (ii) must be similarly sensitive. Different types of moral assessment are made on different grounds. That one set of grounds makes outcomes relevant to one type of assessment, does not imply that the same is true of other grounds and other assessments, even if those other assessments are related and their grounds overlap. So it does not follow from the fact that D2 committed an additional wrong—which D1 did not—that D2 is more blameworthy.

The engagement view offers an explanation of why she is not. To assess blameworthiness, on that view, we must look at the situation *ex ante*—up to the moment D behaves as she does. We must ask whether there is some aspect of D's practical reasoning that reflects badly on D, and we must ask how bad the reflection is in light of how D reasoned. To assess wrongdoing, in contrast, we must look at the situation *ex post*—we must look, *inter alia*, at the results of D's actions. Only then will we know whether D failed to conform to reasons not to succeed. Only then will we know which wrong D committed.

We can see the implications of this analysis by considering a further example:²⁹

Doctors: D3 and D4 are doctors who give V different drugs in order to alleviate severe pain. Each takes reasonable care to ensure that the drug in question is suitable for V. Unfortunately, the drug administered by D3 causes V to have an unusual allergic reaction. V dies as a result.

Neither D3 nor D4 is blameworthy for what they do. Why not? Because an agent's guiding reasons are not necessarily the same as *the reasons by which an agent should be guided*. When we look from the *ex ante* perspective, there is no deficiency to be found in the practical reasoning of either doctor. Indeed, we should hope that both would

²⁹ The example is drawn from Simester (n 24) 202-203.

act the same way in future. Having taken reasonable care to check that the proposed drug was safe, both engaged with the applicable guiding reasons in a manner that is beyond criticism. This is not to say that there is no moral difference between the two. That difference becomes clear when we switch to the *ex post* perspective. D3 had decisive reason not to kill V. Yet this—tragically—is what she did. Even those who are reluctant to call this act wrong, can accept that D3 has reasons for regret that D4 does not. It is worse for her to live with such reasons than to live without them.³⁰ What does not follow is that we should think worse of D3 because of what she did—that it is appropriate to view *her* more negatively in respect of her actions. Judgments of this kind, the engagement view tells us, are to be made *ex ante*. From that perspective—looking at D3’s practical reasoning itself—we find nothing less than what we should reasonably expect.

4. Back to Criminal Law: A Third Way

Our interest is in what the criminal law would look like if it reflected the blameworthiness of those who commit wrongs. Such a criminal law, we assumed, would conform to (Blaming): it would blame D for the wrong she committed, in proportion to her blameworthiness for committing it. We have claimed that in *Bowmen*, D1 and D2 commit different wrongs: D2 killed V; D1 did not. To blame D1 and D2 accurately for the wrongs they committed, the criminal law must convict D1

³⁰ For argument along these lines, see Gardner (n 1) 54-61.

and D2 of different crimes.³¹ So we should accept what we earlier called (C-S): in a criminal law that accurately reflected blameworthiness, some crimes would be defined with reference to outcomes. We should side with Moore and Gardner on criminalisation.

Though they commit different wrongs, we have claimed that D1 and D2 are equally blameworthy for committing them. This is entailed by what we called the engagement view. To blame D1 and D2 in proportion to their blameworthiness requires that they be punished to the same degree. The point generalises. Where the practical reasoning of two agents is deficient in the same way, the two agents are equally blameworthy irrespective of how things turn out – irrespective, that is, of the outcomes of their actions. A criminal law that conformed to (Blaming) would punish them to the same degree. This is to accept what we earlier called (P-I). It is to side with Ashworth on the quantum of punishment.

There are significant differences between our explanation of (P-I) and that offered by Ashworth. Recall that Ashworth's explanation relied on a claim about preconditions. Because control is a precondition of blameworthiness, and outcomes lie outwith our control, the outcomes of our actions cannot make us more blameworthy. As Moore pointed out, this control condition saddles one with a *reductio*. It is important to see that the explanation we have offered is not subject to this objection. We did not argue that outcomes are unable to increase

³¹ Implicit in this claim is that the difference between these two wrongs is sufficiently important to warrant the creation of distinct crimes. More on that in section 5.

blameworthiness because of a control condition, or because of any other precondition. Instead, we argued that this inability drops out of an appropriate account of the grounds of blameworthiness—it follows not from an account of eligibility for blame, but from an account of what makes us blameworthy.

To see the difference this makes, think again about our circumstances. As we already mentioned, many aspects of our circumstances are beyond our control, and many of those aspects bear on whether we try. Had V not walked through the field, D1 and D2 would not have tried to kill V. The control condition implies that D1 and D2 are not blameworthy for trying. The view we have defended has no such implications. The fact that D1 and D2 do not control whether V walks through the field does nothing to show that there is no deficiency in their practical reasoning. So it does not establish that they are not blameworthy for what they do. But their practical reasoning surely *does* reflect badly on them, and this makes it the case that they are blameworthy for their wrongs. To be clear, this does not mean that our circumstances can *never* bear on the degree to which we are blameworthy. That they can explains the existence of excusatory defences like duress. A claim of duress asserts that, in choosing to do the wrong thing, the defendant exhibited no deficiency of practical wisdom *given the circumstances she faced*. Indeed, granted the environment in which she acted, she lived up to what we can reasonably expect of her.³² On the engagement view, it is this kind of argument that must be made to

³² See Gardner (n 17) ch 6.

show that circumstances bear on D's blameworthiness. The mere fact that D does not control her circumstances is neither here nor there.

None of this is to say that blameworthiness has no preconditions. In fact, our view implies that it must. Recall that to assess whether someone is blameworthy is to make a type (ii) assessment. It is to assess whether a particular kind of *connection* exists between an agent and an action. There can be no such connection unless we are dealing with a moral agent—with someone whose actions are responsive to a set of values with which she identifies.³³ This is, however, a much weaker precondition than the condition of control. Imagine that D is kidnapped by T and comes firmly to endorse the values of her captor. Years later, she robs V for the reason that this is what her newfound values tell her she should do. All else being equal, there is no doubt that we are dealing with a moral agent. True, that agent had no control over the person she became. But this does not show, it seems to us, that D is ineligible to be blamed for robbery. If we are right, blaming judgements are to that extent shallow. They require that we be moral agents, but they do not require us to be responsible for being the moral agent we are.

The last few paragraphs all help make the following simple point. Having proposed a view that does not rely on the control condition, we have proposed a view that cannot be reduced to absurdity by that condition's implications. The first *reductio* does not apply. And the same is true of the second. According to the second

³³ Or at least, does not reject. There are deep waters here, but for the purposes of this paper we need not swim in them.

reductio, our moral agency cannot be restricted solely to our endeavours,³⁴ because if there are no moral reasons (not) to succeed then, at least in standard cases, there are no moral reasons (not) to try either. We do not claim that our moral agency *is* so restricted. In fact, we have claimed that it is not—that there are moral reasons (not) to succeed as well as moral reasons (not) to try, and that success sometimes changes which wrong one has committed. This is why we should accept (C-S). What we have denied is that this increases the blameworthiness of those who succeed. This is why (P-I) should be preferred to (P-S). There is nothing in either of these claims that threatens (even in standard cases) to shrink our moral agency to vanishing point. So the Honoré/Gardner *reductio* is also inapplicable.

Before moving on, it is worth noting something the engagement view does *not* imply. It does not imply that we can identify the degree to someone is blameworthy without reference to the kind of wrong they committed. Imagine D5 lies in order to obtain some benefit. Imagine D6 kills V for the very same reason. We have more powerful guiding reasons not to kill than not to lie. So the deficiency in D6's practical reasoning is greater than the deficiency in D5's. On the engagement view, D6 is more blameworthy than D5. Gardner writes that 'the starting point of the blameworthiness inquiry' is 'the action that was wrongful'.³⁵ We agree. The action that was wrong is the very action for which D should be blamed. And the guiding reasons which make D's action wrong play a key role in the examination of whether D's practical reasoning was defective, and how defective that reasoning was. That

³⁴ *Mutatis mutandis* for our endangerments.

³⁵ Gardner (n 17) 231.

we should begin with D's wrong, however, is no reason to reject the engagement view.

5. Defining Crimes: Outcomes and Fair Labelling

Any attempt to combine (C-S) with (P-I) might be thought to be open to a number of objections. We consider two such objections across this section and the next. In the process, we offer some further reasons to accept the account of blameworthiness sketched in section 3.

Both objections begin with the following observation. We have claimed that the facts that make acts wrong, and the facts that ground blameworthiness, are different facts. What make an act wrong are the guiding reasons that bear on that act. What makes us blameworthy is our engagement with those guiding reasons. That engagement must be, in some respect, deficient. We have also assumed that blaming responses should have certain *objects*—we have assumed throughout that those who act wrongly should be blamed *for* their wrongs. This is one aspect of the norm we have called (Blaming).

The first objection has it that if we are right about the grounds of blameworthiness, we must be mistaken about the proper objects of blaming responses. If our engagement with guiding reasons makes us blameworthy, it is for deficient engagement that we should be blamed. By our own admission, the

existence of such a deficiency, and its gravity, are unaffected by outcomes. So a criminal law that accurately reflected blameworthiness would define all crimes in the inchoate mode. We are back, in other words, to (C-I).

In meeting this objection, let us assume that we are dealing with blaming responses that, like those of the criminal law, publicly communicate blame. D1 and D2 are equally blameworthy. Are there reasons for our blaming responses to communicate that D1 and D2 committed wrongs of different kinds? Interestingly, it is in Ashworth's own work that we find the foundations of a positive answer. According to the so-called principle of fair labelling, 'widely felt distinctions between kinds of offences and degrees of wrongdoing' – at least where the feelings in question are reasonable – should be 'respected and signalled by the law'.³⁶ There is, of course, a widely-felt distinction between the wrong of killing and the wrong of merely trying to kill. We have argued that the distinction is not only reasonable but accurate. Those who kill fail to conform to reasons not to succeed that mere attempters do not. To fairly label, our offence-definitions should mark that distinction.³⁷

³⁶ A Ashworth, *Principles of Criminal Law*, 6th edn (OUP 2009) 78.

³⁷ One might object that if we include outcomes in the definition of offences, this will inevitably be taken to convey that, in the eyes of the law, they make D more blameworthy. We disagree. Particularly if bad outcomes do not alter the degree to which D is punished, their inclusion in offence-definitions need not be taken to indicate additional blameworthiness.

One might, of course, ask why the principle of fair labelling is worthy of adherence.³⁸ Consider two explanations. The first appeals to the interests of defendants. Because murder is widely felt to be a more serious wrong than theft, thieves have an interest in not being publicly blamed for murder. For the very same reason, attempted murderers have an interest in not being publicly blamed for murder.³⁹ One might object that this gives us no reason to include outcomes in the definition of crimes. If we exclude outcomes, attempted murderers and murderers will be publicly blamed for the same wrong. Both will be blamed for attempted murder. Yet no-one has a valid complaint about being publicly blamed for the less serious of two wrongs she has committed. So there is no defendant who has a valid complaint if outcomes are excluded from the definition of crimes. In response to this objection, let us grant that murderers have no valid complaint about being blamed for attempted murder. It remains the case that *attempted* murderers can validly complain. If murderers and attempted murderers are blamed for attempted murder, the message sent by the criminal law is that anyone convicted of attempted murder may well have been proved to have committed murder. Murder is – rightly – widely thought to be the more serious wrong.⁴⁰ Attempted murderers have an interest in criminal laws that do not send this message about what they did. The same is true of

³⁸ For general discussion, see J Chalmers and F Leverick, 'Fair Labelling in Criminal Law' (2008) 71(2) *Modern Law Review* 217.

³⁹ Do attempted murderers have this interest only if a murder conviction conveys that they are more blameworthy? The answer is no. There is no reason to think that assessments of type (ii) are the only moral assessments of significance. To falsely portray someone as a wrongdoer is itself objectionable, whether or not one also portrays that person as blameworthy (this is why it is worse to be portrayed as excused, when one is in fact justified, even though both portrayals preclude blameworthiness). We all have an interest in others' refraining from portrayals of this kind.

⁴⁰ We have powerful reasons not to kill, in virtue of the interests others have in remaining among the living. Only those who kill fail to conform to these reasons. This is why murder is the more serious wrong.

all defendants who commit inchoate crimes. Because these defendants *do* have an interest in having different labels applied to those who merely try, versus those who also succeed, the principle of fair labelling supports the inclusion of outcomes in offence-definitions.

A second reason to care about fair labelling appeals to the interests of victims. It is plausible to think that, when allocating labels, we should not only concern ourselves with what the criminal law says about D. We should also concern ourselves with what it says about what D *did to V*.⁴¹ One reason to think this is given by Victor Tadros. In his words:

‘Public solidarity in a political community can be developed only if each person is confident that her moral significance and her interests are adequately reflected in the scheme of laws. Beyond her confidence that each person will obey the law, each person must recognize that the scheme of laws is one whose content demonstrates that the community is properly committed to respecting her as an autonomous agent, and to ensuring that her interests are protected and advanced through common adherence to the law.’⁴²

⁴¹ Some crimes are, of course, victimless. For those crimes, only the first reason discussed here will apply.

⁴² V Tadros, ‘Fair Labelling and Social Solidarity’ in L Zedner and JV Roberts (eds), *Principles and Values in Criminal Law and Criminal Justice* (OUP 2012). For related thoughts on the significance of outcomes to conviction, see A Cornford, ‘Resultant Luck and Criminal Liability’ in RA Duff et al., *The Structures of Criminal Law* (OUP 2011).

To show solidarity with V, the law must send the message that V's interests matter to the political community. We all have interests in not being harmed, that extend beyond our interests in having others not try to harm us. Our prospects in life are affected by the former in ways they need not be affected by the latter. It follows that we have additional interests in not being victims of harmful wrongs. One way to show that these interests are taken into account by the political community is for legal verdicts to register the fact that D's wrong harmed V. For the courts to declare merely that D tried to harm V – when V was actually killed, maimed, or molested – risks sending the message that V's interests in avoiding these outcomes are a matter of indifference.

The point is only partly communitarian. We are owed things by the state not only as citizens but also as human beings. The criminal law is concerned with wrongs committed by and against people within its jurisdiction, even if they are not members of the local political community. Imagine such a wrong is committed against V. She is no part of the community, so the question of solidarity between members does not arise. The very least that justice requires, it is plausible to think, is that the criminal law's verdict acknowledge what was actually done to V – that V was in fact killed, maimed, or molested. Why so? Because criminal verdicts are part of the public record. Their contents are a reflection of what is taken to be significant about what D did. This is why it is not enough for those verdicts to acknowledge that *a* wrong was perpetrated. Doing justice to V – treating her interests in avoiding harmful wrongdoing as morally important – requires a verdict that articulates *the*

way in which V was wronged. Courts deliver such verdicts only if outcomes – like the fact that D killed, maimed or molested V – are included in the definition of crimes.

6. Imposing Punishment: Engagement and the Grounds of Blameworthiness

We turn now to the second objection, which moves in the opposite direction from the first. Embedded in (Blaming) is the assumption that blaming responses are *for* wrongs. According to the second objection, if we are right about the proper objects of blaming responses, we must be mistaken about the grounds of blameworthiness. If (i) we are properly blamed for the wrongs we commit, and if (ii) outcomes sometimes make it the case that we commit additional wrongs, then (iii) those outcomes must also make us more blameworthy. In a criminal law that accurately reflected blameworthiness, these defendants would be punished more. We are back, in other words, to (P-S).

Assume (i) is correct. We have accepted (ii). Success *can* make a moral difference: D2 in *Bowmen* commits a wrong that D1 does not. A rather stipulative version of the second objection has it that we can move straight from that acceptance to (iii): if D2 should be blamed for more, D2 must also be more blameworthy. This is not so.⁴³ We should distinguish between that for which we are properly blamed, and that which makes us blameworthy – between the proper objects of blame, and the grounds of blameworthiness. It is the latter that determine how blameworthy we

⁴³ As noted by Thomson and by Zimmerman. See the articles cited in n 18.

are. Recall that to be blameworthy for ϕ ing is for ϕ ing to reflect badly on us. The worse the reflection, the more blameworthy D is for having ϕ ed. The grounds of blameworthiness determine whether ϕ ing so reflects, and how bad the reflection is. So they determine the degree to which we are blameworthy.

A more sophisticated version of the objection focuses on the grounds of blameworthiness themselves. According to Moore, blameworthiness has two independent grounds. Imagine that D ϕ s, and that it is wrong for D to do so. If D's practical reasoning is deficient—if her engagement reflects badly on her—this is both necessary and sufficient to make D blameworthy for having ϕ ed. But if D's ϕ ing causes harm, this makes D more blameworthy than D otherwise would be.⁴⁴ So outcomes also belong among the grounds of blameworthiness. The engagement view captures only part of the truth.

In defending this rival about grounds Moore relies partly on the *reductio* we considered in section 2. We already explained why that *reductio* is inapplicable here. But Moore also writes that '[t]he argument that has convinced me in the past, and convinces me still, is an experiential argument'.⁴⁵ That argument runs as follows:

- (1) We feel greater guilt when we wrongly cause harm than when we do not;
- (2) This feeling is virtuous—it is what decent people feel;
- (3) Virtuous emotions are epistemically reliable in the judgments they produce;

⁴⁴ Moore (n 1) 191-193.

⁴⁵ Moore (n 17) 129.

- (4) The judgment produced by feelings of guilt is that we are blameworthy;
- (5) Therefore, the fact that we feel greater guilt when we wrongly cause harm is reason to believe that we are more blameworthy for doing so.

Even if (1) is correct, why accept (2)? Presumably, what makes an emotion virtuous is that it is appropriate to feel that emotion in the circumstances at hand. One might plausibly argue that it is appropriate to *feel* more guilty when we *are* more guilty – when we are more blameworthy for what we have done. But as Gardner points out this makes the experiential argument viciously circular.⁴⁶ The argument is meant to give us reason to think that outcomes make us more blameworthy. It turns out, however, that the truth of premise (2) depends on our having already accepted that this is the case.

In response, Moore argues that it is sometimes appropriate to feel guilty when we are *not* blameworthy. He mentions cases of necessity, in which D's actions are justified because they are the lesser evil.⁴⁷ For Moore, it is appropriate that D feel guilty about what D did. But it is far from clear that such feelings *are* appropriate. What is appropriate, it seems to us, is the emotion Bernard Williams called agent-regret: the regret one feels in virtue of the fact that one's actions failed to conform to reasons.⁴⁸ When we cause harm to others it is appropriate that we feel additional regret because we failed to conform to additional reasons: our reasons not to succeed

⁴⁶ Gardner, 'Wrongdoing by Results: Moore's Experiential Argument' (n 13).

⁴⁷ M Moore, 'Four Friendly Critics: A Response' (2012) 18 Legal Theory 491, 532.

⁴⁸ B Williams, *Moral Luck* (CUP 1981).

in causing others harm. Be that as it may. Even if Moore is correct, the examples do not help his argument. For they suggest that it is appropriate to feel guilty not only when we *are* blameworthy, but also when we are *not* because our conduct is morally justified. And this suggests that virtuous feelings of guilt, far from being a reliable indicator of blameworthiness, are actually a rather poor indicator of the same. They may indicate that our conduct was blameworthy. But they may equally indicate that it was blameless. If feelings of guilt may or may not be indicative of blameworthiness, the greater guilt we feel when we cause harm supplies little reason to believe that this makes us more blameworthy. The experiential argument, it seems to us, is a weak one.

It does not follow, of course, that we should accept the engagement view. We still need to know what can be said in its favour. Consider the following case:

Satanic Cult: Members of a gang kidnap someone. They strap him to a chair behind a partition, with the barrel of a rifle running through a small hole and pointing at the kidnapped victim's heart. The rifle holds twenty rounds, and the gang loads it with nineteen blanks and one live shell. No one has any idea which shell is the live one. Twenty gang initiates who want to become full-fledged gang members are required as a condition of membership to pull the rifle's trigger once. (The initiates are unknown to one another and are not in any way acting in concert.) Each initiate pulls the trigger, and at the

conclusion of the rite, the victim is found to have been killed. No one can tell which initiate fired the round that killed him.⁴⁹

Larry Alexander and Kim Ferzan think all are equally blameworthy. They write:

‘Each person knew there was a chance that he might kill an innocent victim and chose to take that risk to join the cult. Each has shown the same disrespect for the victim’s life. ... The one who pulled the trigger that fired the live round, whoever it was, is no more culpable than the other nineteen.’⁵⁰

We agree. The same would hold, we would add, if each of the initiates had actually wanted to kill the victim. The explanation of all this, it seems to us, is given by the engagement view. Because the practical reasoning of each initiate is deficient in the very same way, they are blameworthy for their wrongs to the same degree. Because their practical reasoning is unaffected by what happens next, their degree of blameworthiness is unaffected by the outcomes of what they do.⁵¹

This is not to claim that the initiates in *Satanic Cult* are morally on a par.

Antony Duff writes that those who did not kill the victim have reason to feel ‘relief

⁴⁹ See Alexander and Ferzan (n 6) 175.

⁵⁰ Alexander and Ferzan (n 6) 175-176.

⁵¹ Some things Alexander and Ferzan say suggest that they accept this explanation. They suggest, for instance, that a ‘practical reasoning account’ can escape Moore’s *reductio*. But they also write that [u]ltimately, our position rests on the assumption that the control we have over our choices—our willings—is immune to luck’ (ibid 189-190). Like Ashworth, they appeal to the control condition. But because uncontrolled factors do bear on whether we choose, it is plain that our choices are *not* immune to luck. As we pointed out in section 4, our argument makes no such appeal.

at the non-completion of a wrong', and that this 'properly conditions' the criminal law's response to their actions.⁵² Once again, we agree. Those who commit different wrongs should, all else being equal, be convicted of different crimes. Those who kill commit a different wrong to those who merely try to do so. They commit a different wrong because the reasons not to succeed in killing are reasons to which killers alone fail to conform. Those initiates who are fortunate enough to conform to these reasons have reason to feel relief at their good fortune. But as we have seen at pains to point out in this article, it does not follow that we should think better of them in respect of their actions. So it does not follow that they are less blameworthy for what they did.

We would add that the engagement view also has intuitive appeal when it comes to *praiseworthiness*. Suppose that Diane intends to help Peter in some charitable endeavour by gifting him a painting, which she thinks is worth \$10,000. The painting turns out to be worth \$100,000, enabling Peter to achieve many more of his charitable goals. Diane's actions have generated additional benefits, and thus a better outcome. But she is no more praiseworthy for the gift she makes. Those additional benefits played no part in her practical reasoning.

Intuitions aside, there is a deeper reason to prefer the engagement view to its rivals. Notice first that when deciding between accounts of the grounds of blameworthiness, we are deciding which aspects of D's action bear on our

⁵² RA Duff, 'Thank God I Failed' in HM Hurd (ed), *Moral Puzzles and Legal Perplexities: Essays on the Influence of Larry Alexander* (Cambridge: Cambridge University Press, 2018) ch 4.

judgments of D with respect to that action, by making it appropriate that we think worse of D in light of those aspects.⁵³ One might well be tempted to include outcomes on the list. But to do so is to accept that we are made worse—and are appropriately judged to be so—by the contributions fate makes to how things turn out. It leaves the moral light in which we are properly viewed at the mercy of passing ducks and errant gusts of wind. As Nagel famously claims, any view that has these implications ‘seems irrational as a moral position’.⁵⁴

Many writers—including Ashworth, Alexander, and Ferzan—assume that this line of thought drives us to the control condition. But like Moore, Nagel sees that this cannot not be right. If we endorse that condition:

‘Eventually nothing remains which can be ascribed to the responsible self, and we are left with nothing but a portion of the larger sequence of events, which can be deplored or celebrated, but not blamed or praised.’⁵⁵

This is not a conclusion many people can easily accept:

‘We are unable to view ourselves simply as portions of the world, and from inside we have a rough idea of the boundary between what is us and what is not, what we do and what happens to us, what is our personality and what is

⁵³ As Nagel puts it, ‘when we blame someone for his actions we are not merely saying it is bad that they happened, or bad that he exists: we are judging *him*, saying he is bad....’ See Nagel (n 8) 25.

⁵⁴ Ibid 31.

⁵⁵ Ibid 37.

an accidental handicap. We apply the same essentially internal conception of the self to others. About ourselves we feel pride, shame, guilt, remorse—and agent-regret. We do not regard our actions and our characters merely as fortunate or unfortunate episodes—though they may also be that. We cannot simply take an external evaluative view of ourselves—of what we most essentially are and what we do. And this remains true even when we have seen that we are not responsible for our own existence, or our nature, or the choices we have to make, or the circumstances that give our acts the consequences they have. Those acts remain ours and we remain ourselves, despite the persuasiveness of the reasons that seem to argue us out of existence.’⁵⁶

Our inability to give up on this ‘essentially internal conception of the self’ is one thing that ultimately compels us to reject the control condition. Another is the way that condition has us conceive of others. Suppose that D assaults a policeman because of a threat that otherwise she will be killed.⁵⁷ The implications of the control condition may initially seem appealing: because D lacks control over the threat, she cannot be blamed for her response to it. An acquittal certainly seems to be the correct result.

Closer inspection, however, reveals a different picture. The control condition has us conclude that because D does not control the threat posed, D is not morally

⁵⁶ Ibid 37.

⁵⁷ See, eg, *DPP v Lynch* [1975] AC 653.

responsible for her actions. If that is right, there is no question of our evaluating *what D does next*: there is no question of our asking whether D's *response* to those circumstances reflects well or badly on D. After all, we have already concluded that D is not eligible for this type of moral assessment. What is crucial here is that it is precisely such an assessment that D would surely wish us to conduct, and which D would invite a criminal court to conduct by pleading a defence of duress. To plead such a defence is not to deny but to *assert* one's eligibility for praise and blame. It is to claim that one lived up to the expectations others reasonably have of one in such circumstances; that one should not be blamed because there is no deficiency of practical reasoning to be found in one's response. If we refuse to assess these claims—as the control condition would have us do—we refuse to view D as the kind of morally responsible agent D understands herself to be. No doubt there are special cases in which this refusal is justified. Infancy is one such case. It is hard to believe, however, that competent adults who act under duress are to be equated with children.

Indeed, the implications of the control condition go further than that. For that condition not only mandates that we refuse to evaluate the responses of agents under duress. Given that aspects of our circumstances always lie beyond our control, it mandates that we refuse to evaluate *anyone's* response to the circumstances in which they find themselves: we must give up on the thought that the moral light in which we appropriately view others is determined (even in part) by how they respond to the challenges their circumstances create. Yet there is something

fundamentally disrespectful, it seems to us, about this way of interacting with others: about ceasing to have expectations of others as practical reasoners who may acquit themselves well (or badly) in the face of whatever challenges life throws at them. That it would have us take this dim view of our fellow human beings is one reason the control condition is so hard to accept.

Imagine we do reject the control condition. It may seem that we are then compelled to accept the position Nagel condemned as irrational – that outcomes make us more blameworthy than we would otherwise be. But for reasons we have already given, this dilemma is merely apparent. There is a third way that lies in between viewing ourselves and others in ways we find unacceptable, and giving a degree of moral significance to outcomes that appears irrational. That third way is given by the engagement view. On that view, passing ducks and errant gusts of wind cannot alter the degree to which we are blameworthy.⁵⁸ Only our practical reasoning can do that. Neither is this an *ad hoc* restriction. Rather, it is the natural thought to have for beings who endorse an ‘essentially internal’ conception of themselves. For if anything belongs on our side of the divide between ourselves and the external world, it is surely our engagement with the reasons that bear on what we do. Grounding our blameworthiness in that engagement is compatible with – indeed supported by – the self-conception we find inescapable. Neither does it require that

⁵⁸ They can, we have claimed, alter which wrong we commit. But as the Gardner/Honoré reductio shows us, this is inescapable. See also Simester, ‘Causation in (Criminal) Law’ 133 *Law Quarterly Review* 422-425.

we take a disrespectful view of our fellow human beings. This is another reason, it seems to us, to prefer the engagement view.

7. Conclusion

Let us return, in closing, to the law. We have argued that in a criminal law that accurately reflected blameworthiness:

(P-I): The quantum of punishment imposed by criminal courts would be unaffected by outcomes.

(C-S): Some crimes would be defined with reference to outcomes.

We defended (C-S) by reference to the principle of fair labelling. The outcomes of our actions, we argued, often alter which wrongs we commit. There are significant, widely-felt distinctions between wrongs that are partly constituted by outcomes—like murder, rape or robbery—and wrongs that are not—like attempted murder, rape, or robbery. We fairly label when our offence-definitions draw those distinctions.

We defended (P-I) by reference to what we called the engagement view. On that view, it is the character of our practical reasoning that makes us blameworthy for our wrongs, and which determines how blameworthy we are for their commission.

The outcomes of our actions come too late to alter the character of that reasoning. So they cannot alter the degree to which we are blameworthy for what we do.

One payoff of this second defence is that it explains the criminal law's doctrine of concurrence: the requirement that the *actus reus* and *mens rea* of a crime must coincide in time.⁵⁹ Strictly speaking, only the *behavioural* and *circumstantial* elements of an *actus reus* must occur concurrently with D's *mens rea*. If D is driving and accidentally hits a cyclist, the fact that she realises the cyclist is V, her enemy, and rejoices while he is dying of his injuries will not make her guilty of a homicide offence.⁶⁰ Conversely, if D deliberately poisons V, and V takes some hours to die, the fact that D repents in the meantime will not absolve her of murder.⁶¹

One might ask, though, why the criminal law does not apply its concurrence requirement to the *consequence* elements of the *actus reus*, either in addition to or instead of the element of behaviour. The answer is given by the engagement view. The wrong D commits is partly constituted by the consequence proscribed by criminal law. But D's blameworthiness with respect to that consequence is fixed when (and by the manner in which) the guiding reasons not to bring it about figure in D's practical reasoning. We need to know about D's *mens rea* at *that* moment—the moment of her behaviour—in order to evaluate D's blameworthiness for the wrong she committed. We have assumed throughout that criminal conviction

⁵⁹ *Fowler v Padget* (1798) 7 Term Rep 509, 101 ER 1103.

⁶⁰ We assume that V's injuries are such that death is inevitable.

⁶¹ *Jakeman* (1982) 76 Cr App R 223.

communicates blame, and that D should be blamed for wrongdoing only if – and to the extent that – she is blameworthy for her wrong. The concurrence requirement, as traditionally interpreted, helps the criminal law to satisfy this condition.

Our focus has been on what the criminal law would look like if it accurately reflected D's blameworthiness. One might point out that in the real world, accurate reflection is all but impossible. Might a criminal law concerned with tracking blameworthiness use outcomes as a useful proxy? Might it, for instance, punish those who cause harm more than those who do not, on the basis that this will tend to result in those who are more blameworthy being punished to a greater degree? We do not deny that this might be so. Imagine D has decisive reasons not to harm V, and D acts for those reasons. In such a case, D's conduct is guided by the prospect of harm – she takes the production of that harm to count in favour of altering her behaviour, in ways that make the production of harm more likely.⁶² Compare a case in which D takes her reasons not to harm V to count decisively *against*. Here, D takes the prospect of harm to count in favour of alterations of the opposite kind – those which make it less likely that V will be harmed. All else being equal, practical reasoning of the first kind seems more likely to result in harm than practical reasoning of the second: that D's engagement with the applicable reasons is deficient increases the probability that this will be the outcome of what she does.⁶³

⁶² We draw here on T Nagel, 'Agent-Relativity and Deontology' in his *The View From Nowhere* (OUP 1986).

⁶³ This may help explain why many intuitively associate harmful outcomes with greater blameworthiness. Some writers take this association to be evidence that causing harm is among the grounds of blameworthiness. An alternative explanation is that causing harm is evidence of those

For ease of exposition, we have focused primarily upon intended wrongdoing. But the view we have defended has similar implications for acts of endangerment. It is one thing unreasonably to risk harming others. It is another to cause the harm that was risked. To do the second thing is to commit an additional wrong not committed by those whose conduct is fortuitously harmless. If we are right, however, this does not mean that one actor is more blameworthy than the other: 'successful' endangerment is *ceteris paribus* no more blameworthy than its inchoate counterpart.

One might think that this commits us to the conclusion that the law of inchoate offences should be expanded: that as well as a general inchoate offence of attempt, there should be a general inchoate offence of endangerment. Not so. There are good reasons why the law should not criminalise all unreasonable risk-taking. To do so would open up the prospect of more invasive policing, and would necessarily result in selective prosecutions. At the same time, there is a plausible correlation between successful endangerment and blameworthiness: similar to that which, we have suggested, applies in the case of attempts. In practice, the likelihood that D's act of endangerment goes on to cause harm will vary with the manner and dangerousness of D's behaviour. The greater the risk taken, the more probable it is that harm will occur. All else being equal, it is more blameworthy to impose greater risks on others.

A requirement of actual harm may therefore be a useful proxy for

grounds: i.e., evidence of deficient engagement with guiding reasons. If this second explanation is correct, we should expect the intuitive association to cease when it is clear those who cause harm, and those who do not, engage with the applicable reasons in the same way. This, we have suggested, is precisely what happens in cases like *Satanic Cult*.

blameworthiness.⁶⁴ In turn, given the desirability of avoiding overcriminalization, it may be apt for the law to limit endangerment liability to those who actually cause harm.

These conjectures, which are no more than that, in no way conflict with the analysis offered here. That outcomes roughly track blameworthiness—if they do—does nothing to suggest that they are among its grounds. In fact, it suggests the opposite. If outcomes are merely a useful proxy for blameworthiness, then they do not themselves make it the case that we are more blameworthy for bringing them about. And if the reason outcomes are a useful proxy is that they reliably indicate deficiencies in D's practical reasoning, this supports the conclusion that it is D's practical reasoning that makes D blameworthy. This is one of the conclusions for which we have argued here.

⁶⁴ Or so it has been suggested. See, eg, N Richards, 'Luck and Desert' (1986) 65 *Mind* 198, 200-202; Wolf (n 27) 7; Enoch and Marmor (n 4) 415-16. The possibility is rejected by Cornford (n 42) 47.