

Culpability Without Ability: A Hard Determinist Proposal

YUYANG MA, *DUKE UNIVERSITY*

This essay defends a coherent form of hard determinism against the charge that, if agents are never able to do otherwise, our practices of moral advice and blame lose their ground. I formulate this pressure as the Culpability Dilemma. On the one hand, powerful incompatibilist arguments deny any genuine ability to do otherwise, whether determinism or indeterminism is true; on the other hand, ordinary moral life still trades in judgments like ‘you should have...’ and ‘it’s your fault...’, judgments that seem to presuppose exactly that ability. To get out of this dilemma, I challenge van Inwagen’s Moral Advice Argument by proposing that second-person ‘ought’ claims are not primarily reports of an antecedent normative fact that presupposes present categorical ability but are practical interventions aimed at influencing what an agent will do. On a deterministic picture, correct advice need only satisfy a weaker condition: if the advice were given and taken up, it would causally issue in the advised action. I then revisit Broad’s ‘ought implies can’ by distinguishing comparative, deliberative, and retrospective uses of ‘ought’. These revisions allow the hard determinist to preserve the intelligibility of advice and much of the grammar of fault and blame while rejecting the claim that culpability requires the ability to do otherwise. The essential move is the reinterpretation of Broad’s retrospective “ought”: showing that past-tense moral verdicts reduce to counterfactual claims about correct prior guidance and therefore do not presuppose categorical ability.

1. Introduction

We often blame people for certain states of affairs. We say a careless driver is at fault for an accident, or that a dishonest politician is to blame for a scandal. When we make such judgments, a trivial underlying assumption seems to be that: there must be an occasion on which the person we blame was able to do otherwise. We can more formally capture this thought down in the following Culpability Thesis:

If a bad state of affairs S is truly the fault of a person x , then there is some earlier time t and some course of action A such that:

- (1) At t , x was choosing between two or more incompatible courses of action, one of which is A ;
- (2) at t , x ought to perform A instead of what x actually did;
- (3) if x had performed A , S would not have occurred;
- (4) and at t , x was able to perform A .

The Culpability Thesis interacts in a crucial way with the following three familiar philosophical theses:

Ability Thesis: On at least some occasions, when an agent is deciding between two or more incompatible courses of action, the agent is able to perform each of them.

Thesis of Determinism: The past and the laws of nature determine a unique future.

Thesis of Indeterminism: The past and the laws of nature do not determine a unique future.

Both determinists and indeterminists can offer powerful arguments that their respective views rule out the kind of ability posited by the Ability Thesis. For the hard determinist (one who holds that the Thesis of Determinism is true and that it is incompatible with the Ability Thesis), Peter van Inwagen's Consequence Argument purports to show that if the past and the laws of nature jointly determine every future event, then no agent is ever able to do otherwise than what she actually does; the ability to do otherwise would require either changing the past or violating a law of nature, neither of which is within anyone's power. For the hard indeterminist (one who holds that the Thesis of Indeterminism is true and that it is incompatible with the Ability Thesis), the Mind Argument powerfully shows that indeterminism fares no better: if an agent's action is not determined by her prior mental states, then whether she performs it is a matter of luck or chance, and an action that occurs by chance is not one the agent has the ability to perform or to refrain from performing in the relevant sense.¹ Given that determinism and indeterminism are mutually exclusive and jointly exhaustive, if neither is compatible with the Ability Thesis, it follows that the Ability Thesis is false. Yet the problem remains that in ordinary life, we continue to speak and judge as if people are sometimes at fault. When we think about Nazi officers who oversaw mass killings, we say: "These crimes were their fault" or "they ought not to have done this," which, together with the Culpability Thesis, appears to require exactly the sort of ability that both sides have just ruled out.

C. D. Broad found himself in a similar situation in his inaugural lecture at Cambridge University, *Determinism, Indeterminism, and Libertarianism*, where he argued a libertarian view (one that affirms the Ability Thesis) is hard to maintain given the incompatibilist arguments. At this point, people like Broad, who deny the Ability Thesis on the strength of these incompatibilist arguments face a dilemma with two primary paths forward: (i) defend that "nothing is ever anyone's fault" is not merely counterintuitive but actually true, and then explain the nature of this powerful illusion that sustains ordinary fault and blame attributions or (ii) argue that the Culpability Thesis is false, i.e., moral responsibility does not require any ability to do otherwise. This dilemma is what I will call the Culpability Dilemma. Broad still believes the Culpability Thesis to be self-evident. His reasoning is that whenever we make a retrospective judgment of obligation ("you ought not to have done X"), we universally admit as relevant the defense "I could not help doing X" or "I could not possibly have done Y." The fact that this is treated as a legitimate rebuttal shows, for Broad, that there is some sense of "could" in which "ought" and "ought not" entail "could"; an action is, in his words, "obligable if and only if it is, in a certain sense, substitutable."² So (ii) is not really an option for him.

Peter van Inwagen also builds his own case to defend the Culpability Thesis and block path (ii) in *Free Will and Moral Responsibility: Once More Unto the Breach*, by connecting the concepts of fault, moral advice, and the ability to act, in what I shall call the Moral Advice Argument.³ His argument is built on the core claim that: if "you ought to do X" is the correct answer to the question "What should I do?", then the agent is able to do X.

But does the Moral Advice Argument successfully establish the link between culpability and ability or "ought" and "can", such that one cannot coherently give advice or make fault and blame attributions if they reject the Ability Thesis? This is the question I wish to explore in what follows. Van Inwagen's core claim about advice, as I will show, is a specific application of the Broadian principle: that "ought" and "ought not" can apply to an action only if the agent could have done otherwise; so to establish culpability without ability requires showing that the retrospective "ought" (the "ought" of blame and moral verdict) does not commit us to the categorical sense that Broad takes to be essential.

My strategy proceeds in stages. I will first (§2) offer a way around van Inwagen's Moral Advice Argument by reconceiving advice as a causal intervention, which weakens the link between correct advice and present categorical ability. I will then (§3) turn to the deeper Broadian principle itself and argue that once we distinguish the comparative, deliberative, and retrospective uses of "ought," the retrospective "ought" can be reinterpreted as a counterfactual claim about correct prior guidance rather than as a marker of cate-

¹Peter van Inwagen, *An Essay on Free Will*, rev. ed. (Oxford: Oxford University Press, 1986); and Peter van Inwagen, *Thinking about Free Will* (Cambridge: Cambridge University Press, 2017), chap. 7, 90–110, <https://doi.org/10.1017/9781316711101>.

²C. D. Broad, "Determinism, Indeterminism, and Libertarianism," in *Ethics and the History of Philosophy: Selected Essays* (London: Routledge, 2014), 197, <https://doi.org/10.4324/9781315823058>.

³Peter van Inwagen, "Free Will and Moral Responsibility: Once More unto the Breach," *Belgrade Philosophical Annual* 38, no. 1 (2025): 13, <https://doi.org/10.5937/BPA2501007I>.

gorical ability. This reinterpretation is the essential move to show that the hard determinist can retain the full grammar of “you ought to have done otherwise” without accepting Broad’s claim that such judgments presuppose the ability to have done otherwise. Building on these two strands, I will (§4) demonstrate how the hard determinist can make sense of fault and blame ascriptions and dissolve the Culpability Dilemma via path (ii).

2. The Moral Advice Argument

Van Inwagen’s Moral Advice Argument seeks to demonstrate that assigning fault for a state of affairs logically implies that the agent in question was, at some point, able to act otherwise. The core of the argument can be formulated in three steps. Step one is a simple and reasonable principle of fault, which states that:

“If a certain state of affairs is someone’s fault, there is something that person did not do and should have done (or did and should not have done).”⁴

For example, if it’s Henry’s fault that the cat died, then Henry should have fed the cat, or he should not have performed a cat-disabling experiment (but he did). The object of blame is typically a state of affairs (the cat’s death), which is traced back to a failure to meet an obligation (a thing Henry should have done). Step two is a trivially true proposition:

“If an agent is deciding between *A* and *B*, and they ought to do *A*, then the correct answer to the question ‘what should I do?’ is ‘you ought to do *A*.’”⁵

This step simply blocks the objection that the “ought” of advice might differ in content from the moral “ought”: if there is a fact of the matter about what one should do, then stating that fact constitutes correct advice. When an advice columnist in “The Ethicist” gives correct advice, the advice is correct precisely because it aligns with what the person morally ought to do. Step three argues that:

“If ‘you ought to do *x*’ is the correct answer to an agent’s deliberative question ‘what should I do?’, then the agent is able to do *x*.”⁶

This is the critical step that establishes the ‘ought implies can’ principle. This step is, in effect, a particular application of a principle that Broad articulated in his treatment of obligability: that “ought” and “ought not” can only apply to an action that is, in the relevant sense, substitutable—one that the agent could have performed or could have left undone.⁷ Van Inwagen’s Step 3 applies this principle to the specific context of advice. The disagreement in what follows will therefore concern not only van Inwagen’s particular formulation, but the deeper Broadian commitment that stands behind it, which I will elaborate on in §3.

To illustrate why advising someone to perform an action that is beyond their ability is senseless, van Inwagen used the following example:

Suppose Planner is giving his friend advice on where to take Alison for dinner. Planner advises his friend to take Alison to Auberge du Vieux Puits in Paris, despite knowing full well that his friend cannot afford the trip and could not get there in time.

To see how each steps play together, we can apply them to the Henry feeding the cat example again.

1. Suppose the cat’s death is Henry’s fault (starting premise).
2. By step one, Henry should have fed the cat.
3. By step two, if Henry deliberates whether to feed the cat or go to Camp David, the correct answer to “What should I do?” is: You ought to feed the cat.

⁴van Inwagen, “Free Will and Moral Responsibility,” 15.

⁵van Inwagen, “Free Will and Moral Responsibility,” 15.

⁶van Inwagen, “Free Will and Moral Responsibility,” 16.

⁷Broad, “Determinism, Indeterminism, and Libertarianism,” 195-197.

4. By step three, for that advice to be correct, Henry must have been able to feed the cat.
5. Therefore, if the cat's death was Henry's fault, Henry was able to feed the cat.

2.1. Moral Advice as Attempt to Influence

The problem of this argument lies in step three which seems to presuppose a hidden principle:

(CA. α) In order for X to give Y the correct advice, X must believe that Y is able to do the thing advised.

Planner's advice, "you ought to take Alison to Auberge du Vieux Puits for dinner" is a wrong answer to his friend's question "where should I take her to dinner?", because Planner does not believe that his friend is able to do the thing advised, violating CA. α .

If we build CA. α into step three, we get:

Strong Step Three (SS.3): If "You ought to do x " is the correct answer to an agent's deliberative question "What should I do?", then the agent is (at the time of advice) able to do x .

Following SS.3, van Inwagen's point is that anyone who rejects the Ability Thesis simply cannot give out any moral advice that is consistent with their belief in that rejection; to the question "what should I do?", they cannot correctly answer "you ought to do x ." To illustrate his reasoning, suppose that person X is a hard determinist. If Y asks X : "What should I do, A or B ?", from X 's epistemically limited perspective, X knows that Y will only do what Y does; since A and B mutually exclusive, X must believe that Y is not able to do at least one of them, but X does not know which. Therefore, on that ground, X would have to say to Y : "I cannot give you any advice, because I don't know what you're able to do."

However, this is not necessarily the case. In this scenario, what X could say is that: "once I gave Y the advice, then, Y would be able to do the thing because he would follow my advice, so it's worth a try." To elucidate, it helps to distinguish two senses in which giving advice might be "worth a try." At the metaphysical level, under hard determinism, whether X 's advice will causally influence Y is already settled: the prior state of the world and the laws of nature determine whether the advice will be taken up, just as they determine everything else. In this sense, it is not "worth a try"—the outcome is fixed. But at the epistemic level, X does not know what Y is determined to do. X does not know whether the advice will be part of the causal chain that leads Y to act, or whether it will fall on deaf ears. From X 's epistemically limited perspective, giving advice is worth a try in the sense that, for all X knows, the advice might be the very causal factor that determines Y 's action.

This reframes what X can say when Y asks, "What should I do, A or B ?" X need not claim to know that Y is presently able to do either. Rather, X can reason as follows: "I do not know, at this moment, whether Y will be able to do what I advise. But I do know that, if I give the advice and it is taken up, it will causally contribute to Y 's doing the advised action." Since X cannot know in advance whether the advice will be efficacious, the rational course is to advise as if it might be—to treat the giving of advice as a causal intervention whose success is epistemically open, even if metaphysically settled.

The "able to do A " condition need not hold at the time of asking but can be satisfied at the time of action after the advice. In this sense, we can now replace CA. α with a more nuanced principle:

(CA. β) In order for X to give Y the correct advice, X has to believe that Y is able to do the thing advised after receiving the advice.

Going back to step three, notice that the strong version secretly builds in the requirement of the agent's ability at the moment of advice. It states you may give it as correct advice only if the agent, independently of your giving it, already has the power to comply. A hard determinist can reject this as the wrong picture of advice. In a deterministic picture, when Y comes to X and asks, "What should I do, A or B ?", X is not standing outside the causal order trying to read off, like a fortune teller, whose actions are already in Y 's ability set. X is trying to influence Y 's decision by adding a cause or a determining factor, namely, the advice.

Under this conception of advice, a hard determinist can then build CA. β into step three and have a weaker version:

Weak Step Three (WS.3): If “You ought to do x ” is the correct answer to an agent’s deliberative question “What should I do?”, then the agent is (after the advice) able to do x when the time for action comes.

We can see how the revised WS.3 can change the Moral Advice Argument by applying it to the Henry’s cat case. Steps 1–3 remains the same, but the revised steps 4 and 5 now reads:

4. If “You ought to feed the cat” is the correct answer, then Henry must have been able to feed the cat (at the time of action) as a direct result of the advice itself.
5. Therefore, if the cat’s death was Henry’s fault, then there is a correct piece of advice such that, had it been given, Henry would have been able to feed the cat at the time of action.

The cat’s death being Henry’s fault only proves that the correct advice “feed the cat” would have been a causally efficacious factor in bringing about that action. It does not prove that Henry had any ability to choose that path independently of being caused to do so by the advice.

Those are different modal claims. SS.3 licenses an inference from “correct advice” to alternative possibility: holding fixed the actual past and the laws, Henry had both options genuinely open. WS.3 only licenses an inference from “correct advice” to counterfactual, advice-enabled possibility: in a nearby possible world where the past includes the advice, Henry feeds the cat. A natural objection is: what if the advice fails? What if Henry receives the advice and, being determined by his prior psychological states and circumstances, does not take it up? This is a real possibility, and the hard determinist should concede it since WS.3 does not guarantee that advice will always be efficacious. What it guarantees is weaker: that the correctness of advice does not presuppose any ability on the agent’s part at the time of asking. Whether the advice succeeds is a further empirical question, one whose answer is metaphysically determined but epistemically open to the adviser. The hard determinist gives advice under the same epistemic conditions as anyone else—not knowing for certain whether it will work, but having reason to think it might, and recognizing that the giving of advice is itself a potential causal factor in the outcome.

2.2. Moral Advice and Non-Cognitivism

The account of advice I just sketched relies on a simple thought: when someone asks, “What should I do, A or B ?”, they are not merely asking for a description of how the moral universe already stands; they are asking for guidance that is meant to influence what they will do.

This way of thinking sits comfortably with many forms of non-cognitive or moral subjective views of morality. Suppose a person thinks that there are moral truths about things, then moral advice is just supposed to be a report about what the moral universe says. Asking, “what should I do, A or B ?” is like asking, “which number is bigger, 7 to the 3 or 3 to the 7?” There’s an answer to this question, and answering it is to report the truth, not to guide someone’s arithmetic. That’s the point of cognitivist morality, that my account of moral advice is inconsistent with.

This non-cognitive aspect is a useful note to register explicitly how my conception of advice differs from van Inwagen’s. In his view, second-person ‘ought’-statements are primarily informative or reportorial. When an agent asks, “What should I do, A or B ?”, van Inwagen understands this as a request for a judgment about an objective normative fact; the adviser’s role is to provide an answer that accurately reports how things stand: what the right act is. What the advisee then does with that information is, strictly speaking, up to them and lies outside the content of the advice itself. In this picture, the adviser does not “have a dog in the fight” between the agent’s competing inclinations; the adviser is not trying to strengthen one desire or create a new one, but to tell the truth about which option is right. By contrast, the hard-determinist reinterpretation treats second-person “ought”-claims as practical interventions in a causal process. When X says to Y , “ABC is the best restaurant,” the primary role of this utterance is not to neutrally report a normative fact for Y to do with as Y pleases, but to strengthen or create an inclination in Y and thereby

help determine what Y will in fact do. The content of the utterance is inseparable from its role in trying to steer behavior down a certain path. This is not to say that the advice lacks evaluable content: X can still be wrong about which restaurant is best, and if X is wrong, then X has given bad advice. The point is that the function of the utterance in the deliberative context is interventionist rather than reportorial. Whether the advice is correct (by whatever evaluative standards one endorses) and whether it is causally efficacious are two separate questions. The hard determinist can accept that advice is assessable for correctness while maintaining that its role in the practice of advising is to intervene in a causal process whose outcome is metaphysically determined but epistemically open to both adviser and advisee.

The Moral Advice Argument implicitly presupposes the more “objective-report” conception. My proposal is that a hard determinist can instead adopt the “interventionist” conception, and once that shift is made, the move from “correct advice” to present categorical ability in SS.3 no longer follows.

2.3. Moral Advice and Hard Indeterminism

Another important note is that this conception of moral advice does not work for the hard indeterminists and only works under a specific conception of hard determinism.

To see why, it helps to state more precisely what $CA.\beta$ requires. On the closest-world semantics for counterfactuals, $CA.\beta$ is satisfied when the following nested counterfactual is true:

Had the agent been advised to do A , then, had the agent tried to do A , the agent would have done A .

The outer counterfactual shifts us to the nearest possible world in which the agent receives the advice; the inner counterfactual, evaluated at that world, asks whether trying to comply would have resulted in compliance. Crucially, this is a claim about what counterfactuals are in fact true—a matter of the modal structure of the world—and not a claim about what X knows. X gives advice under epistemic uncertainty about whether these counterfactuals hold, just as a doctor prescribes a treatment without certainty that it will work. What matters is whether the counterfactuals are true, not whether X can verify them.

With this in hand, we can see why the same proposal fails for the hard indeterminist. Suppose X is a hard indeterminist. If Y asks X , “What should I do, A or B ?”, X does not know, prior to giving advice, what Y is able to do. But crucially, even after giving the advice, X cannot be sure that Y will be able to do the action advised—and this is not merely an epistemic limitation. In X ’s picture of the world, the relevant counterfactual may lack a determinate truth-value: even in the closest possible world in which Y receives and tries to follow the advice, indeterministic factors may prevent compliance. The problem is the metaphysical level, not epistemic. Under determinism, if the advice enters the causal history, the closest world in which Y receives and tries to follow it is one in which Y succeeds (given causal stability). Under indeterminism, no such assurance is available even at the level of counterfactual truth, because chance may intervene between trying and succeeding.

This contrast also reveals a tacit presupposition in the deterministic case. My account of advice assumes a deterministic world with relatively low modal fragility at the level of ordinary deliberation: small variations in how advice is delivered do not produce wildly divergent outcomes within the timeframes relevant to deliberation. This is a claim about the modal structure of the world, not about what the adviser knows. It does not require infallible modal knowledge on X ’s part; it requires only that, as a matter of fact, the closest worlds in which the same advice is given with minor variations (delivered in 19 seconds rather than 20 or phrased slightly differently) are worlds in which the same outcome obtains. If the world were so chaotic that this were not the case, then the relevant counterfactuals would be false, and $CA.\beta$ would not be satisfied—but this would be a failure of the world to cooperate, not a failure of X ’s knowledge.

My argument therefore supposes that, within the timeframes and levels of description relevant to ordinary deliberation, advice functions as a reasonably stable causal influence. This supposition goes beyond standard hard determinism, but it is precisely the sort of stability that our everyday advisory practices already take for granted.

3. Does ‘Ought’ Imply ‘Can’?

So far, I have focused on van Inwagen’s Moral Advice Argument and how a hard determinist can reinterpret the link between advice and ability. The deeper question remains: how can the hard determinist coherently use the language of ‘ought’ at all, or in other words, does all ‘ought’ imply ‘can’? Having loosened van Inwagen’s particular formulation, in this section, I turn to Broad’s treatment of this deeper question in his inaugural lecture, *Determinism, Indeterminism, and Libertarianism*. I will argue that Broad’s conception of ‘obligability’ is tied to a particular retrospective use of ‘ought’, and he does not clearly distinguish this from other, forward-looking uses of ‘ought’ as advice. On that basis, I reconstruct his two senses of ‘ought’ in a way that a hard determinist can retain.

3.1. Broad on ‘Ought Implies Can’

Broad starts from a familiar kind of retrospective moral judgment about a past action. We say things like: “You ought not to have done *A*; you ought instead to have done *B*.” He then notes that if the person wants to resist this judgment, they have two options. Either they say: (a) “I could have done *A* instead of *B*, but you are wrong about what I ought to have done; in fact, *A* was what I ought to have done,” or they say: (b) “I could not help doing *A*,” or “I could not possibly have done *B*.” In (a), the agent grants that ‘ought’ or ‘ought not’ applies, but disputes which action is the one they ought to have done. In (b), the agent denies that ‘ought’ or ‘ought not’ applies at all, because he or she was not able to do otherwise. Broad takes (b) to rest on something like a shared background assumption: if an agent literally was not able to do otherwise, then neither ‘ought’ nor ‘ought not’ applies to that action. This is what leads him to the thesis that ‘obligability entails substitutability’.

In Broad’s terminology, an action is “obligable” if it is the kind of thing to which ‘ought’ or ‘ought not’ can be applied. An action is “substitutable” if it was done but was able to have been left undone, or it was left undone but was able to have been done.⁸ He then tries to make this more precise by clarifying that the kind of substitutability entailed by obligability is a ‘categorical substitutability’, one that does not share the logical form: “I would have been able to do otherwise if I had willed otherwise” or “I would have willed otherwise if I had willed differently in the past.” Broad’s view can hence be summarized in the following thesis:

“‘Ought’ or ‘ought not’ can only apply to an action if the action was done but was able to have been left undone, or it was left undone but was able to have been done.”⁹

Broad also holds that there are two general senses of ‘ought’. The first he attributes to Spinoza and calls the ‘comparative ought’; This ‘ought’ can be applied to both humans and non-humans alike which simply states that one state of affairs ranks higher than its competitors. For example, suppose a person says, “This soup ought to have less salt.”; what the person meant is that its flavor is worse than that of a properly seasoned soup. Similarly, when applied to human conduct, “You ought not to cheat at poker” may mean that cheating falls below the standard of an average decent person. Such claims are entirely independent of whether the agent is able to meet that standard; this sense, when applied to humans, is perfectly consistent with an absence of categorical substitutability or the denial of the Ability Thesis.

The second sense of ‘ought’ is what Broad called, the ‘categorical ought’. When a person feels guilty and thinks, “I ought not to have done that,” the person is not just saying, “my character is worse than average”; the person is implying that he or she was able to have and should have chosen differently, independent of his or her pre-existing character. This is the sense of ‘ought’ that implies categorical ability and hence making the Culpability Thesis more believable.

⁸Broad, “Determinism, Indeterminism, and Libertarianism.” In the actual paper, Broad used the classical phrase “could have” as opposed to “was able to”. However, given that “could have” is ambiguous as mentioned in 1.1, I assume what he meant by “could have” is really just “was able to”, as opposed to “might have”.

⁹Broad, “Determinism, Indeterminism, and Libertarianism,” 197.

3.2. The Grammars of ‘Ought’

One important note for Broad’s arguments is that he never distinguishes between the retrospective ‘ought’ and the moral advice ‘ought’. In fact, Broad’s examples and formulations in the lecture almost all concern past-tense judgments:

“You ought not to have done *X*.”,

“You ought to have done *Y* instead.”,

“He ought not to have let himself become a slave to morphine,” etc.

For Broad, the grammatical form here is that we look back at what someone did (or failed to do) and pronounce a verdict. Broad calls these retrospective judgments of obligation about past actions, and his whole discussion of ‘obligability’ and ‘substitutability’ is framed in terms of such judgments.

What he does not do is to distinguish this retrospective ‘ought’ from other very common uses of the word we discussed in §2, namely: the future-directed ‘ought’ of moral advice (“You ought to tell the truth tomorrow”). Instead, Broad slides from his careful discussion of past-tense ‘ought’ into a more general-sounding principle of ‘ought implies can’. It is then natural for readers to take this as a constraint on all uses of ‘ought’.

But this slide can be resisted. We can distinguish between different roles that ‘ought’ plays in our language and re-interpret this retrospective ‘ought’ in a way that aligns with the future-directed ‘ought’.

Drawing on the earlier discussion of moral advice, we can now distinguish at least three useful senses or roles of ‘ought’:

- (i) First, we still have the ‘comparative ought’. This Spinozan sense of ‘ought’, as sketched out by Broad, is certainly compatible with hard determinism.
- (ii) Second, we have the ‘deliberative ought’. This is the sense that appears when *X* asks, “What should I do, *A* or *B*?” and *Y* answers, “You ought to do *A*.” Given the picture developed in §2, this ‘ought’ is tied to the practice of advice as an attempt to influence or guide action and hence can also be used consistently under hard determinism.
- (iii) Third, we have the ‘retrospective ought’. This is the past-tense ‘ought’ that Broad focuses on when one looks back and makes a normative verdict about a past action. Based on Broad’s view, this is a ‘categorical ought’ that is rendered false if the agent was not able to do otherwise.

The problem is how can hard determinists interpret the ‘retrospective ought’ in a way that does not commit them to the categorical sense. We can use the insight from WS.3 in §2; a hard determinist can analyze retrospective judgments of the form “*X* ought (not) to have done Φ ” by reducing them to counterfactual claims about correct prior advice. More precisely:

- (R1) “*X* ought (not) to have done Φ at *t*” $\Leftrightarrow$ “If, at some suitable time *t**, prior to *t*, *X* had asked ‘What should I do?’, then the correct answer would have been: ‘You ought (not) to do Φ .’”

For a hard determinist, to make a retrospective judgment and say “you ought to have done or ought not to have done something” is like saying, “the advice I would have given you if you’d asked beforehand is *x*.” On this analysis, a retrospective ‘ought’ is a claim about what the correct deliberative ‘ought’ would have been, had *X* sought advice or guidance beforehand. The truth-conditions of the retrospective “You ought to have done Φ ” are therefore identical to the truth-conditions of the corresponding advisory claim “You ought to do Φ ” evaluated at an earlier time *t**. In this way, the retrospective ‘ought’ inherits the interpretation of the deliberative ‘ought’ the hard determinist adopts in CA.β, without introducing any further commitment to categorical ability.

Therefore, when *X* says to *Y*, “you ought to have done *A* or ought not to have done *B*,” *X* might be doing at least two things:

- (1) Evaluating *Y*'s past action relative to a standard ("Y's action falls below the standard for a decent person"), which is a comparative 'ought',
or
- (2) Indicating what the correct advice would have been and, simultaneously, issuing a forward-looking recommendation ("In similar situations, you ought to do *A* and not do *B*"), which is effectively a deliberative 'ought'.

Under this interpretation, there are still essentially two senses of 'ought', both of which are compatible with hard determinism, undermining the 'ought implies can' principle as argued by Broad.

4. Out of the Culpability Dilemma

So far, I have argued that a hard determinist can make sense of second-person moral advice by treating "You ought to do *A*" as a causal intervention rather than as a neutral report about an independent normative fact. A natural next question is whether the same strategy can be extended to fault and blame.

The hard determinist can answer this by extending the same causal interventionist account from advice to blame. To say "It is your fault that *S* happened" has at least two components. First, there is a descriptive, cognitive component: it attributes to the agent a certain causal and normative relation to *S*; it notes that an action or omission by the agent played a role in the deterministic history leading to *S* and evaluate that action as falling below some accepted moral standard. In this sense, the hard determinist can say "You are at fault," to mean "Your behavior was one of the causes of *S* and manifested a defect relative to the standard we endorse." No appeal to categorical abilities is required for that much.

Secondly, there is a forward-looking, non-cognitive component. In ordinary life, we use "You are to blame" to produce feelings of guilt and remorse to shape the agent's future disposition or to reinforce a shared norm in the community. In the interventionist picture, ascriptions of fault and blame are themselves attempts to influence the future causal chain in much the way that advice is.

This allows the hard determinist to reasonably reject the ability requirement found in clause (3) of the Culpability Thesis and instead substitute a weaker condition, parallel to the Weak Step Three from §2.1. Accordingly, we can construct a hard deterministic reformulation of the Culpability Thesis as follows:

A bad state of affairs *S* is truly the fault of a person *x* if and only if (1*) *x*'s action or omission was a part of the causal history of *S*, (2*) that action or omission falls below a standard that we are entitled to enforce, and (3*) there is some correct pattern of guidance or advice such that, had *x* been appropriately exposed to and responsive to it earlier, *x* would have behaved differently and *S* would not have occurred.

On this view, fault-ascriptions, like advice, are part of the machinery by which a deterministic system regulates itself.

We can now see how this returns us to the Culpability Dilemma. Given that the practice of any fault and blame ascriptions require only causal connection to bad outcomes, deviation from standards we are prepared to enforce, and a forward-looking role for blame and advice as instruments for influencing behavior, we can consistently deny 'ought implies can' under hard determinism, and path (ii) remains open to get out of the Culpability Dilemma.

Bibliography

- Broad, C. D. "Determinism, Indeterminism, and Libertarianism." *Ethics and the History of Philosophy: Selected Essays*. Routledge, 2014. <https://doi.org/10.4324/9781315823058>.
- van Inwagen, Peter. *An Essay on Free Will*. Rev. ed. Oxford: Oxford University Press, 1986.
- van Inwagen, Peter. *Thinking about Free Will*. Cambridge: Cambridge University Press, 2017. <https://doi.org/10.1017/9781316711101>.
- van Inwagen, Peter. "Free Will and Moral Responsibility: Once More unto the Breach." *Belgrade Philosophical Annual* 38, no. 1 (2025): 7–19. <https://doi.org/10.5937/BPA25010071>.
- van Inwagen, Peter. "Does 'Ought' Imply 'Can'?" *Free Will: Historical and Analytic Perspectives*, edited by Marco Hausmann and Jörg Noller, 313–33. Cham: Springer, 2021.