

---

---

# THE JOURNAL OF PHILOSOPHY

VOLUME CXX, NO. 10, OCTOBER 2023

---

---

## THE DEONTIC PRIMACY OF ACTIONS?\*

**E**thical theories must answer two questions. What are the actions we ought to perform? This is the Extensional Question. Why ought we to perform these actions? This is the Explanatory Question.<sup>1</sup> There is a good case to be made that, with the right machinery, for any plausible theory, we can devise a counterpart that delivers the same extensional answer.<sup>2</sup> In determining which theory is best, it is a theory's explanatory answer that matters most.<sup>3</sup>

We can categorize the various answers to the Explanatory Question based on whether they hold that the oughts governing actions are explained by the oughts governing non-actions. Kantians, for example,

\*I would like to thank Winnie Sung, Douglas Portmore, and Aldrin Relador for helpful comments on earlier drafts. I would also like to thank several anonymous reviewers and editors for the careful attention they gave the argument. This research was supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (RG62/19 (NS)).

<sup>1</sup>I borrow the extensional/explanatory labels from Benjamin Kiesewetter, "Are All Practical Reasons Based on Value?," in Russ Shafer-Landau, ed., *Oxford Studies in Metaethics*, vol. 17 (Oxford: Oxford University Press, 2022), pp. 27–53.

<sup>2</sup>Many early Utilitarians at the extensional level were Divine Command Theorists at the explanatory level—for example, William Paley. And more recently it has been argued that we can consequentialize or deontologize any plausible theory. See, for example, Jennie Louise, "Relativity of Value and the Consequentialist Umbrella," *The Philosophical Quarterly*, LIV, 217 (2004): 518–36; Jamie Dreier, "In Defense of Consequentializing," in Mark Timmons, ed., *Oxford Studies in Normative Ethics*, vol. 1 (Oxford: Oxford University Press, 2011), pp. 97–118; Douglas W. Portmore, "Consequentializing," *Philosophy Compass*, IV, 2 (2009): 329–47; Douglas W. Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011); and Paul Hurley, "Consequentializing and Deontologizing: Clogging the Consequentialist Vacuum," in Mark Timmons, ed., *Oxford Studies in Normative Ethics*, vol. 3 (Oxford: Oxford University Press, 2013), pp. 123–53. For criticism, see Campbell Brown, "Consequentialize This," *Ethics*, CXXI, 4 (2011): 749–71.

<sup>3</sup>For defense of this idea, see Franz Dietrich and Christian List, "What Matters and How It Matters: A Choice-Theoretic Representation of Moral Theories," *The Philosophical Review*, CXXVI, 4 (2017): 421–79; and Marius Baumann, "Consequentializing and Underdetermination," *Australasian Journal of Philosophy*, xcvi, 3 (2019): 511–27.

hold that, if you ought to perform some act, this is explained by the fact that you ought to will the corresponding maxim.<sup>4</sup> Why ought you to refrain from telling this false promise? Because so acting would be in accordance with a maxim that you ought not will. The deontic properties of an agent's actions are, according to the Kantian, explained by the deontic properties of her willings. By contrast, Mooreans hold that the deontic properties of actions are not explained by the deontic properties possessed by non-actions. The deontic properties of actions are, according to Mooreans, explained by the evaluative ranking of the outcomes. If you ought to perform some action, this is explained by how the value of its outcome compares to the value of the outcomes of the available alternatives. Why ought you to refrain from telling this false promise? Because telling it would bring about a worse outcome than the outcome of one of your other available alternatives. In terms of the Explanatory Question, the Kantian rejects, while the Moorean accepts, the

DEONTIC PRIMACY OF ACTIONS: It is not the case that the deontic properties of actions are explained by the deontic properties of non-actions.

In what follows, I will lay out an argument against the Deontic Primacy of Actions. This argument aims to show that the deontic properties possessed by an agent's actions are explained by the deontic properties possessed by her attitudes.<sup>5</sup>

<sup>4</sup>I should stress that my main aim is to give a feel of the sort of position one might adopt in rejecting the deontic primacy of actions. Perhaps this is not Kant's view. But it is a natural way of interpreting several of his claims concerning the Formula of Universal Law. Michael Cholbi, for example, seems to interpret Kant in this way. See his *Understanding Kant's Ethics* (Cambridge, UK: Cambridge University Press, 2016), section 3.1.

<sup>5</sup>Allen Gibbard, T. M. Scanlon, and Michael Smith seem to endorse this claim. But they do not offer much by way of argument. See Allen Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, MA: Harvard University Press, 1990); T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), section 1.2; Michael Smith, "The Ideal of Orthonomous Action, or the How and Why of Buck-Passing," in David Bakhurst, Brad Hooker, and Marget Olivia Little, eds., *Thinking about Reasons: Themes from the Philosophy of Jonathan Dancy* (Oxford: Oxford University Press, 2013), pp. 51–74; and Michael Smith, "Parfit's Mistaken Meta-ethics," in Peter Singer, ed., *Does Anything Really Matter? Essays on Parfit on Objectivity* (Oxford: Oxford University Press, 2017), pp. 99–120. Recently, Conor McHugh and Jonathan Way put forward two arguments for the claim that reasons for action are fundamentally reasons for intention. See their "All Reasons Are Fundamentally for Attitudes," *Journal of Ethics and Social Philosophy*, xxi, 2 (2022): 151–74. Interestingly, like the argument below, one of McHugh and Way's arguments relies on the Response Constraint. Moreover, though we get to this claim using different arguments, we agree that normative support for actions entails normative support for certain attitudes. McHugh and Way argue that the best explanation of this entailment is that the normative properties of

Put roughly, this argument runs as follows: Any of an agent's actions that possess deontic properties could not obtain without being caused by attitudes that possess deontic properties. When this causal relation holds, the deontic status of an action and the deontic status of the attitudes causally responsible for it cannot have independent explanations. So the deontic properties of the attitudes explain the deontic properties of the actions, the deontic properties of the actions explain the deontic properties of the attitudes, or some third factor explains both. The latter two are non-viable. It is not the case that the deontic properties of the actions explain the deontic properties of the attitudes, nor is it the case that some third factor explains both. Hence the deontic properties of attitudes explain the deontic properties of actions.

I should state upfront that the defense of each premise in this argument will be rather limited. What I hope is of value is seeing the result of putting these ideas together. My main goal, then, is not to offer a sustained defense of each claim. Instead I aim to showcase how a number of plausible claims at the intersection of normative ethics, moral psychology, and the philosophy of action entail the conclusion that what an agent ought to do is explained by the attitudes she ought to have.

Before turning to the argument, two preliminary remarks are in order. First, when I use 'explained by', I mean the everyday sense that is used when answering a why-question—the sense that can be replaced with 'because', 'grounded in', or 'in virtue of'. And one explanation, in this sense, can be open to further explanation. This matters for the coming argument. For instance, some consequentialists—for example, Sidgwick, Smith, and Portmore—claim that the deontic properties of actions are *proximally* explained in terms of the evaluative properties of their outcomes, and then go on to claim that the evaluative properties of these outcomes are *ultimately* explained by the normative properties of the agent's attitudes.<sup>6</sup> Since the Explanatory Question is concerned with the ultimate (not proximal) explanation of the deontic properties of actions, we should say that these consequential-

---

actions are explained by the normative properties of attitudes (specifically intention). The argument here, by contrast, aims to use this entailment, along with several other claims, to show that the normative properties of actions must be explained by the normative properties of attitudes. Since McHugh and Way's arguments are different but compatible with mine, our arguments could be brought together, making for a strong case against the deontic primacy of actions.

<sup>6</sup> See Henry Sidgwick, *The Methods of Ethics*, 7th ed. (London: Macmillan, 1907); Michael Smith, "Two Kinds of Consequentialism," *Philosophical Issues*, xix, 1 (2009): 257–72; and Portmore, *Commonsense Consequentialism*, *op. cit.*

ists end up on the same side as the Kantians.<sup>7</sup> But to say this, we need to claim that those divided over the Deontic Primacy of Actions are divided over whether the deontic properties of actions are *ultimately* explained by the deontic properties of non-actions. Below, to avoid cumbersome locutions, I will drop this qualification. Yet it should be remembered that whenever I speak of the Deontic Primacy of Actions, I intend it to be read as having incorporated the ultimate/proximate distinction.<sup>8</sup>

Now for the second preliminary. I will use 'the evaluative' and 'the normative' to mark the contrast that others mark with 'the good' and 'the right'. Accordingly, by 'evaluative properties' I mean those that can be ascribed with predicates such as 'is good', 'is bad', or 'is valuable'. I will treat deontic properties and reason properties as falling within the normative domain. By 'deontic properties' I mean those properties that can be ascribed with predicates capturing all things considered normative assessments like 'is required', 'is permissible', or 'is optional'. By 'reason properties' I mean those properties that can be ascribed with predicates capturing normative contributions like 'is a reason for', 'is a reason against', or 'is a decisive reason'. I assume there to be a tight connection between deontic properties and reason properties.<sup>9</sup> This assumption bears on how the Deontic Primacy of Actions is to be understood. Given the connection between reasons and oughts, I assume, for example, that if a Kantian held that the deontic properties of actions are explained by the reason properties of willings, this would still qualify as rejecting the Deontic Primacy of Actions. And if a Moorean held that the deontic properties of actions are explained in terms of reasons for actions and reasons for actions are explained in terms of the value of states of affairs, this would still qualify as accepting it. So the central question of this article might be more precisely formulated as follows: Are the deontic properties of actions explained by the normative properties—the deontic properties or reason properties—of non-actions? However, to keep the formulations simple, I will continue to phrase the main claims of the coming argument mostly in terms of deontic properties.

<sup>7</sup> Moore's insistence that goodness cannot be analyzed blocks this route. Once he appeals to goodness to explain the deontic status of actions, he is committed to goodness serving as the ultimate explanation.

<sup>8</sup> I thank an anonymous referee for pressing me to make explicit that what is at issue in this debate is the ultimate (not proximate) explanation.

<sup>9</sup> My use of 'reason' and 'ought' could, for example, be converted into what Roderick M. Chisholm calls "indefeasible requirements" and "defeasible requirements" in his "Practical Reason and the Logic of Requirement," in Joseph Raz, ed., *Practical Reasoning* (Oxford: Oxford University Press, 1978), pp. 118–27.

## I. JUDGMENT-SENSITIVE ATTITUDES ARE NECESSARY

The argument against the Deontic Primacy of Actions begins with the claim that

JUDGMENT-SENSITIVE ATTITUDES ARE NECESSARY: For any agent, *S*, and action, *A*, if it ought to be that *S*'s *A*-ing obtains,<sup>10</sup> then *S*'s having certain judgment-sensitive attitudes is causally necessary for *S*'s *A*-ing to obtain.

This claim holds that, if you ought to perform some action, then certain attitudes are causally required for its performance. For example, suppose that your waving to your neighbor possesses deontic properties. From the fact that this act possesses deontic properties we are licensed to infer something about its causal path if you indeed end up waving to your neighbor: it involves judgment-sensitive attitudes—like beliefs, desires, and intentions.<sup>11</sup>

In calling these attitudes 'judgment-sensitive' I am following the popular terminology introduced by Scanlon.<sup>12</sup> But the notion I have in mind is not importantly different from what Portmore calls "reasons-responsive attitudes"<sup>13</sup> or what Hieronymi calls "commitment-constituted attitudes."<sup>14</sup> What is important for the coming argument is that these attitudes can be under your control. And when these attitudes are under your control, this is sufficient for their possessing deontic properties.<sup>15</sup> So, for example, though non-voluntary, you can have control over your beliefs, because you can reason from

<sup>10</sup>Throughout I use the more cumbersome 'it ought to be that *S*'s *A*-ing obtains' rather than '*S* ought to *A*', because it matches my preferred way of talking about causal necessity. Talking in this way thus affords a uniformity of expression for the various principles defended below. I should stress that nothing I say in what follows hinges on a commitment to oughts relating agents to propositions rather than actions. For more on this debate, see Mark Schroeder, "Ought, Agents, and Actions," *The Philosophical Review*, cxx, 1 (2011): 1–41. I thank an anonymous reviewer for suggesting that this be made explicit.

<sup>11</sup>The list of judgment-sensitive attitudes is not limited to beliefs, desires, and intentions. The list also includes admiration, fear, guilt, indignation, respect, regret, pity, shame, and the like. Among the attitudes on this list, desire is controversial. But I am convinced by the various arguments offered for its inclusion by Warren Quinn, *Morality and Action* (Cambridge, UK: Cambridge University Press, 1994); Scanlon, *What We Owe to Each Other*, *op. cit.*; Judith Jarvis Thomson, *Normativity* (Chicago: Open Court, 2008); Derek Parfit, *On What Matters*, vol. I (Oxford: Oxford University Press, 2011); M. Smith, "Parfit's Mistaken Meta-ethics," *op. cit.*; and Douglas W. Portmore, *Opting for the Best: Oughts and Options* (Oxford: Oxford University Press, 2019).

<sup>12</sup>Scanlon, *What We Owe to Each Other*, *op. cit.*, p. 18.

<sup>13</sup>Portmore, *Opting for the Best*, *op. cit.*, pp. 75–76.

<sup>14</sup>Pamela Hieronymi, "Controlling Attitudes," *Pacific Philosophical Quarterly*, LXXXVII, 1 (2006): 45–74; and Pamela Hieronymi, "Responsibility for Believing," *Synthese*, CLXI, 3 (2008): 357–73.

<sup>15</sup>For further discussion and defense of this claim, see Philip Pettit and Michael Smith, "Freedom in Belief and Desire," this JOURNAL, xcii, 9 (1996): 429–49; Michael

your evidence and your beliefs are responsive to this reasoning.<sup>16</sup> This is why it is not wrongheaded, using the normative sense of ought, to assert that, “You ought to believe that the earth is round.” This makes judgment-sensitive attitudes unique. Even if you just woke up after a long night’s sleep, you cannot respond to this fact as a reason to cease being tired. This is why it is simply wrongheaded, using the normative sense of ought, to assert that, “You ought not to be tired now.”

Below, I will defend two claims that together (along with some plausible assumptions) entail Judgment-Sensitive Attitudes Are Necessary. First, an action of yours possesses deontic properties only if you can perform this action for a reason directly. Second, you perform an action for a reason directly only if your judgment-sensitive attitudes cause you to perform this act directly.

Are there certain features that your actions need to have in order to be up for normative assessment?<sup>17</sup> Common sense answers Yes. But common sense is undecided about which features in particular. Being free, volitional, or intentional are frequently invoked. Yet such answers are controversial and the notions they employ are hard to pin down. In this section, I aim to defend something less controversial and more precise:

*Only Actions Possibly Done for Reasons:* For any agent, *S*, and action, *A*, *S*’s *A*-ing possesses deontic properties only if *S* can *S* directly for a reason.

---

Smith, “Rational Capacities, or: How to Distinguish Recklessness, Weakness, and Compulsion,” in Sarah Stroud and Christine Tappolet, eds., *Weakness of Will and Practical Irrationality* (Oxford: Oxford University Press, 2003), pp. 17–38; Pamela Hieronymi, “The Wrong Kind of Reason,” this JOURNAL, cii, 9 (2005): 437–57; Hieronymi, “Controlling Attitudes,” *op. cit.*; Hieronymi, “Responsibility for Believing,” *op. cit.*; Angela M. Smith, “Responsibility for Attitudes: Activity and Passivity in Mental Life,” *Ethics*, cxv, 2 (2005): 236–71; Angela M. Smith, “Attitudes, Tracing, and Control,” *Journal of Applied Philosophy*, xxii, 2 (2015): 115–32; Miriam McCormick, “Taking Control of Belief,” *Philosophical Explorations*, xiv, 2 (2011): 169–83; Miriam McCormick, *Believing against the Evidence: Agency and the Ethics of Belief* (New York: Routledge, 2014); Ralph Wedgwood, *The Value of Rationality* (Oxford: Oxford University Press, 2017); Conor McHugh, “Epistemic Responsibility and Doxastic Agency,” *Philosophical Issues*, xxiii (2013): 132–57; Conor McHugh, “Attitudinal Control,” *Synthese*, cxciv, 8 (2017): 2745–62; and Portmore, *Opting for the Best*, *op. cit.* There are, of course, different accounts of attitudinal control. But I will simply assume that whatever the correct sense of control turns out to be such that our attitudes are up for deontic assessment, it is the same sense that allows our attitudes when under our control to rationalize our actions. With this assumption, we do not need to defend a particular account of control, which would be a project in its own right.

<sup>16</sup> McHugh makes similar observations in his “Attitudinal Control,” *op. cit.*, p. 2756.

<sup>17</sup> As far as I know, Sidgwick was the first to offer a detailed examination of this question. See his *The Methods of Ethics*, *op. cit.*, Book I, chapter 5, section 2 and Book III, chapter 1, section 2.

Why think an action possesses deontic properties only if it is possible to perform that action for a reason? The answer appeals to the thought that an agent ought to perform some action only if she has the ability to act on the basis of the facts—the normative reasons—generating the ought. A normative reason for an action must be such that the agent can respond to it. This idea is known as the

*Response Constraint:* For any agent, *S*, action, *A*, and fact, *F*, *F* is a normative reason for *S* to *A* only if *S* can *A* directly on the basis of *F*.<sup>18,19</sup>

For a fact to count in favor of an agent acting in a certain way, it must be possible for it to be the fact for which she acts directly. I use ‘directly’ to rule out the possibility that the agent performs some other actions (or forms some attitudes) which in turn causes her to perform the action.<sup>20</sup> Assuming that it is impossible to possess deontic properties in the absence of normative reasons, the Response Constraint tells us that actions that cannot be done directly for reasons cannot possess normative properties. Thus the constraint delivers Only Actions Possibly Done for Reasons.

The Response Constraint is popular because it offers an attractive way of making the cut between the evaluative and the normative. To see this, consider

**Avalanche:** An avalanche is headed down a mountain. If it goes left, it will plow into some people. If it goes right, it will slow to a stop in an empty field. It ought to be that it turns right.

**Driver:** You are driving down a mountain and your brakes go out. If you turn left, you will plow into some people. If you turn right, you will slow to a stop in an empty field. It ought to be that you turn right.

<sup>18</sup> This formulation is modified from Nishi Shah, “A New Argument for Evidentialism,” *The Philosophical Quarterly*, lvi, 225 (2006): 481–98, at p. 484. For others who endorse and defend versions of this principle, see Thomas Kelly, “The Rationality of Belief and Some Other Propositional Attitudes,” *Philosophical Studies*, cx, 2 (2002): 163–96; Niko Kolodny, “Why Be Rational?,” *Mind*, cxiv, 455 (2005): 509–63; Parfit, *On What Matters*, *op. cit.*, appendix A; and Joseph Raz, *From Normativity to Responsibility* (Oxford: Oxford University Press, 2011).

<sup>19</sup> I adopt the reading of the Response Constraint that treats ‘can’ as referring to a specific ability—the ability to do something in the present situation. This is to be contrasted with a reading of ‘can’ as referring to a general ability—the ability to do something in a wide range of situations. For example, here and now you have the general ability and specific ability to read. If tomorrow you will be anesthetized for a short operation, then during that time you will retain the general ability to read but lose the specific ability. For the distinction between general ability and specific ability, see Alfred R. Mele, “Agents’ Abilities,” *Noûs*, xxxvii, 3 (2003): 447–70. For how this distinction can deliver different interpretations of the Response Constraint, see Jonathan Way and Daniel Whiting, “Reasons and Guidance (Or, Surprise Parties and Ice Cream),” *Analytic Philosophy*, lvii, 3 (2016): 214–35.

<sup>20</sup> I borrow this idea from Portmore, *Opting for the Best*, *op. cit.*, p. 25.

Although superficially similar, the ought-statement that concludes Avalanche is importantly different from the ought-statement that concludes Driver. The ought in Avalanche is evaluative. It tells us, of the ways things might go, which is better. The ought in Driver is normative.<sup>21</sup> It tells us that you are required to make things go a certain way. The Response Constraint explains why this is so. That you can act on the fact that turning right would avoid hurting many is what opens you up to normative assessment. The ought in Driver sticks to you as an agent—a being capable of recognizing and responding to the relevant facts. Since the avalanche cannot recognize and respond to these facts, only evaluative assessment applies.<sup>22</sup>

The Response Constraint is important because it provides a principled way of restricting the set of actions that possess deontic properties. We have located a feature that the members of this set have in common: the agent can do them for reasons directly. And this feature brings other features in its wake—features that bear directly on the Deontic Primacy of Actions.

We can now turn to the claim that all acts done for reasons directly are caused by judgment-sensitive attitudes. More precisely, the claim I aim to defend is

*Attitudes as Causal Precursors:* For any agent, *S*, and action, *A*, *S* As directly for a reason only if *S*'s judgment-sensitive attitudes cause *S* to *A* directly.

Attitudes as Causal Precursors claims that if an agent's judgment-sensitive attitudes do not cause her to perform an action, then she does not perform that action for a reason.

<sup>21</sup> This assumes that certain things are true of you in these circumstances. We thus should say, if we want to be more careful, that the ought in Avalanche is unambiguously evaluative, while the ought in Driver is ambiguous. In Driver, it is unclear whether the ought is normative or evaluative. This difference—that the ought is ambiguous in Driver but not in Avalanche—is sufficient to support the Response Constraint. For what we would need to know about you to disambiguate the claim in Driver is precisely what the Response Constraint predicts. We would need to know if you have the specific ability to respond to the reasons in play. I owe this point to an anonymous reviewer.

<sup>22</sup> This paragraph, including the cases and reasoning, follows John Gibbons, *The Norm of Belief* (Oxford: Oxford University Press, 2013), chapter 6; and Portmore, *Opting for the Best*, *op. cit.*, section 1.1.9. Bernard Williams presented a similar line of reasoning earlier in his "Internal Reasons and the Obscurity of Blame," in *Making Sense of Humanity and Other Philosophical Papers, 1982–1993* (Cambridge, UK: Cambridge University Press, 1995), pp. 35–45, at pp. 39–40. I should mention that there are well-known objections to the Response Constraint—for example, surprise party reasons. See Mark Schroeder, *Slaves of the Passions* (Oxford: Oxford University Press, 2007), pp. 33–34. But there are also convincing replies to these objections; for example, see Benjamin Kiesewetter, "You Ought to  $\phi$  Only If You May Believe That You Ought to  $\phi$ ," *The Philosophical Quarterly*, LXVI, 265 (2016): 760–82, at section 3.



We can defend Attitudes as Causal Precursors by showing that actions done for reasons must be actions caused by the agent's judgment-sensitive attitudes. And we can arrive at this conclusion by way of two weaker claims. The first we can call the

*Causal-Psychological Requirement:* For any agent, *S*, and action, *A*, *S*'s *A*-ing is done for a reason only if *S*'s *A*-ing is caused by features of *S*'s psychology.

This requirement does not specify which features in particular need to be doing the causing. That is the job of the second claim, which we can call

*No Control No Rationalization:* For any agent, *S*, and action, *A*, if *S*'s *A*-ing is caused only by that over which *S* lacks control, *S*'s *A*-ing is not done for a reason.

The Causal-Psychological Requirement and No Control No Rationalization taken together lead to the idea that for an agent's action to be done for a reason it must be caused by features of her psychology over which she has control. It is only a short step—one I will assume we should take—from here to Attitudes as Causal Precursors.

The causal-psychological theory of action has been the subject of intense scrutiny.<sup>23</sup> Much of the critical heat concerns wayward causal chains. Yet, since the Causal-Psychological Requirement states only a necessary condition and only applies to actions done for reasons, we can ignore this issue. This requirement is thus very difficult to reject.<sup>24</sup> We are looking for an explanation of a certain sub-class of happenings. To capture the fact that an action was performed, our account needs to explain why something happened. And only causal explanations can explain why things happen. So our account needs to be causal.<sup>25</sup> To capture that the action is done for a reason, our account needs to rationalize the action. We need to pick out items among the causal nexus that could serve as the basis for why the agent acted as

<sup>23</sup> For compelling presentations of the causal theory of action, see Paul M. Pietroski, *Causing Actions* (Oxford: Oxford University Press, 2000); Alfred R. Mele, *Motivation and Agency* (Oxford: Oxford University Press, 2003); and the essays in Jesús Aguilar and Andrei Buckareff, eds., *Causing Human Actions: New Perspectives on the Causal Theory of Action* (Cambridge, MA: MIT Press, 2010).

<sup>24</sup> See, for example, Jerry A. Fodor, *A Theory of Content and Other Essays* (Cambridge, MA: MIT Press, 1990), p. 156.

<sup>25</sup> For further defense, see David Lewis, "Causal Explanation," in *Philosophical Papers*, vol. II (Oxford: Oxford University Press, 1986), pp. 214–40; and Bradford Skow, *Reasons Why* (Oxford: Oxford University Press, 2016), sections 6.1–6.2. For discussion, see Matthieu Queloz, "Davidsonian Causalism and Wittgensteinian Anti-Causalism: A Rapprochement," *Ergo*, v, 6 (2018): 153–72.

she did. And the only plausible candidates are certain psychological states.

Our common ways of describing action done for reasons reflect these facts. We say things like, "She is flipping the switch because it will reset the fuse. Well, actually she's at the wrong fuse box. So, more precisely, she's flipping the switch because she thinks it will reset the fuse." We need these bits of the agent's psychology to be doing the causing of the agent's actions, else we would fail to rationalize them.<sup>26</sup>

The Causal-Psychological Requirement captures the idea that we mark off the class of actions done for reasons by their etiology. They must spring from her psychology for an agent's actions to qualify as having been done for a reason. This requirement is, however, extremely permissive. We could satisfy it with any part of the agent's psychology. No Control No Rationalization aims to restrict the field. If the action is caused by psychological states that are outside of an agent's control, the action fails to qualify as an action done for reasons. To see why this link is on the right track, we can put it to work in the following cases:

Tic: When in awkward social settings, you are overwhelmed by the feeling that you must touch the wrist of your left arm. While at a gathering, you hear the person you are talking to let out an uncomfortable laugh. You start rapidly and repeatedly tapping your left wrist.

Phobia: You have extreme acrophobia. Irrational, recalcitrant fear dominates your psychology at even relatively low heights. Unbeknownst to you, the bottom of the building's elevator that you are riding in is glass. Suddenly, you find yourself staring down at the ground, hundreds of meters below. You cover your face with your hands.

Your tapping your wrist and your covering your face are caused by features of your psychology. Yet these are not actions done for reasons.<sup>27</sup>

<sup>26</sup> For discussion, see Donald Davidson, "Actions, Reasons, and Causes," in *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), pp. 3–16; Bernard Williams, "Internal and External Reasons," in *Moral Luck: Philosophical Papers, 1973–1980* (Cambridge, UK: Cambridge University Press, 1981), pp. 101–13, at p. 102; Kieran Setiya, "Reasons and Causes," *European Journal of Philosophy*, xix, 1 (2011): 129–57; Michael Smith, "Four Objections to the Standard Story of Action (and Four Replies)," *Philosophical Issues*, xxii, 1 (2012): 387–401; and M. Smith, "The Ideal of Orthonomous Action, or the How and Why of Buck-Passing," *op. cit.*

<sup>27</sup> Our linguistic practices are prone to mislead us on this front. For example, if someone asked, "Why were you covering your face like that?," a natural response might be, "Well, the reason is that I have this terrible fear of heights." But this should not be taken as evidence that your covering your face is an action done for reasons. For, as should be clear, this use of reason is merely causal. Citing your phobia does not rationalize your covering your face. G. E. M. Anscombe suggests something in this vicinity

Why? Because they are compulsive—the product of drives or urges outside of your control. They are not the product of psychological states that you have or lack as a result of responding to what you take to be reasons for or against them. They simply befall you. Accordingly, the actions they cause cannot be rationalized in the way characteristic of actions done for reasons.<sup>28</sup>

The Causal-Psychological Requirement holds that all actions done for reasons are caused by psychological states of the agent. No Control No Rationalization holds that no actions done for reasons are caused only by that over which the agent lacks control. Together they take us to the idea that actions done for reasons must be caused by psychological states over which the agent has control. Judgment-sensitive attitudes, in the sense described above, seem to be the only psychological states that fit this description.

It may be objected, however, that the following possibility has not been ruled out: An action possesses deontic properties and so can be done for reasons and so can be caused by judgment-sensitive attitudes, yet the agent ends up performing this act for no reason at all—without any connection to judgment-sensitive attitudes.<sup>29</sup> If this is indeed possible, then Judgment-Sensitive Attitudes Are Necessary is false. But, as I will now argue, this is not possible.

It will help to start with Scanlon. He writes that, “[I]t is the connection with judgment-sensitive attitudes that makes events actions, and hence the kind of things for which reasons can sensibly be asked for and offered at all.”<sup>30</sup> On Scanlon’s view, it is impossible for you to end up performing the very same act for no reason at all, because, without the connection to judgment-sensitive attitudes, you simply do not act. In the absence of these attitudes, there may be your involvement in an event—some movements of your body, say—but there would not be the performance of an action. That, in many cases, seems correct. Take a standard mental state that is not a judgment-sensitive attitude (for example, being dizzy) and turn it into the cause of a bodily movement (for example, falling down). Scanlon’s view rightly holds that

---

when she writes, “[T]hough indeed we readily say, e.g., ‘What was the reason for your starting so violently?’ this is totally unlike ‘What is your reason for excluding so-and-so from your will?’ or ‘What is your reason for sending for a taxi?’” See her *Intention*, 2nd ed. (Cambridge, MA: Harvard University Press, 1963), section 5. For discussion of Anscombe’s view concerning mere mental causes of action, see John Schwenkler, *Anscombe’s Intention: A Guide* (Oxford: Oxford University Press, 2019), section 2.3.

<sup>28</sup>For further argument, see Michael Smith, “The Explanatory Role of Being Rational,” in David Sobel and Steven Wall, eds., *Reasons for Action* (Cambridge, UK: Cambridge University Press, 2009), pp. 58–80.

<sup>29</sup>I thank an anonymous reviewer for pressing me to address this objection.

<sup>30</sup>Scanlon, *What We Owe to Each Other*, *op. cit.*, p. 21.

your falling down on account of being dizzy is not your action; it is something that merely happens to you.

Nevertheless, requiring a connection to judgment-sensitive attitudes might strike some as overly restrictive. Think back to the cases of Tic and Phobia. Since your tapping your wrist and covering your face lack a connection to judgment-sensitive attitudes, they are, on the view we are considering, not actions. And the implication of this local case globalizes. Assuming that judgment-sensitive attitudes are only had by cognitively sophisticated creatures, if we claim that these attitudes are necessary for action, we end up with the view that the only creatures that perform genuine actions are in the same league as mature human beings.

Despite its restrictiveness, Scanlon's view is undeniably tempting. Indeed, I am inclined to assume it is correct. Yet it would be best not to have to come down on such a controversial issue. And we do not have to. Instead of focusing on whether an event is an action in the absence of judgment-sensitive attitudes, we can focus on whether your acting, when done for reasons, could be the same, when disconnected from them. In Tic, for example, however we end up classifying your tapping (action or not) there seems to be an important difference between it and a case of your tapping as a result of your judgment-sensitive attitudes. As a first approximation, your tapping is not the same because in Tic we are unable to attribute the source of the action to something over which you have control.<sup>31</sup> Connected to reasons, an apt way of describing the event suggests possession: You are moving and it is your movement. Disconnected, the suggestion of possession seems out of place: You are moving but it is not your movement.<sup>32</sup> And that seems to be a difference that makes a difference to the very nature of the event.

We can sharpen this idea. Some kinds of thing are what they are in virtue of their causal origin. For such things, nothing can be identical to them unless they share this origin. Even an intrinsically identical burn cannot be the same as a sunburn unless it is also made by the sun. Even an intrinsically identical note cannot be the same as a Ster-

<sup>31</sup> For similar remarks, see Joseph Raz, *Engaging Reason: On the Theory of Value and Action* (Oxford: Oxford University Press, 1999), p. 20.

<sup>32</sup> Here is another place where our linguistic practices might mislead. We might say in Tic that "You are now tapping your wrist" or "Your tapping is annoying." But these expressions can be dismissed as a kind of conversational shorthand. Notice we also say, "You are now digesting your food" or "Your snoring is annoying." And no one who uses these expressions seriously thinks that your food is being digested by you or that the snoring is an action of yours. If digestion and snoring are attributable to someone, it is Mother Nature, not you. For a similar line of thought, see Fred Dretske, "Can Intelligence Be Artificial?," *Philosophical Studies*, lxxi, 2 (1993): 201–16.

ling banknote unless it was also made by the Royal Mint.<sup>33</sup> Even an intrinsically identical ink trail cannot be the same as a genuine Winston Churchill autograph unless it was also made by Winston Churchill.<sup>34</sup> The same holds for your actions when done for reasons. Even an intrinsically identical event cannot be the same as your action unless it was also made by you. An action, if done for reasons, is fundamentally different from an event with no causal connection to reasons at all.<sup>35</sup> Having control over the attitudes that cause these actions makes them yours in a way that sets them apart from other (intrinsically identical) events caused by forces outside your control.<sup>36</sup>

To see why this idea is plausible, consider

**Killing for a Reason:** You wake to find yourself at the edge of a cliff holding a rope from which your enemy hangs. Assessing the situation carefully, you form a belief that he deserves to die, a desire to kill him, and an intention to do so. These attitudes cause you to release your grip. Your enemy dies.

**Phobia Killing:** You have extreme acrophobia. Irrational, recalcitrant fear dominates your psychology at even relatively low heights. You wake to find yourself at the edge of a cliff holding a rope from which your enemy hangs. As you start to assess the situation, you happen to stare down at the ground, hundreds of meters below. You cover your face with your hands, releasing your grip. Your enemy dies.

<sup>33</sup> This example is modified from Fred Dretske, *Perception, Knowledge, and Belief: Selected Essays* (Cambridge, UK: Cambridge University Press, 2000), pp. 209–10.

<sup>34</sup> This example is taken from Daniel Dennett, “Three Kinds of Intentional Psychology,” in Richard Healey, ed., *Reduction, Time and Reality* (Cambridge, UK: Cambridge University Press, 1981), pp. 37–61, at p. 38.

<sup>35</sup> Again, Scanlon’s view that nothing can be an action unless it is connected to judgment-sensitive attitudes would vindicate this idea. So too would certain accounts of event individuation. The view once defended by Donald Davidson—that events are individuated by their causes and effects—would do so. See his “The Individuation of Events,” *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), pp. 163–80, at pp. 179–80. So too would the weaker position held by Peter van Inwagen that “no event could have had causes other than its actual causes.” See his *An Essay on Free Will* (Oxford: Clarendon Press, 1983), p. 170. Yet the same considerations that keep me from appealing to Scanlon’s view apply here as well. It would be best not to saddle the argument with a controversial metaphysical thesis, especially one that is much stronger than what is needed for the argument to go through.

<sup>36</sup> Notice that this claim permits an action of yours, if done for reasons, to be identical with another action so long as the doing of this action has some connection to some reasons. It does not claim that the reasons in play need to be the same reasons. And it does not claim that the action itself needs to be performed for reasons directly. It can, for example, be the side-effect of something else you do for a reason. So, it allows that you do the right thing for the right reasons, do the right thing for the wrong reasons, and do the right thing by doing something else for the right or wrong reasons that in turn causes you to do the right thing. What is ruled out is the possibility that you do the right thing for no reason—absent any causal connection to judgment-sensitive attitudes.

Scanlon's view tells us that in Phobia Killing you do not act. The idea we have been exploring makes a less controversial claim. Even if we assume that there is an action performed in Phobia Killing, it is importantly different from your killing your enemy in Killing for a Reason. In Phobia Killing, you lack control over your fear of heights. That is why "your enemy was killed *by you*" sounds off as a description of what happened in Phobia Killing (but not in Killing for a Reason). It sounds off because there is a real sense in which it was not you doing the killing; it was the phobia doing the work, so to speak.

Your actions done for reasons (as in Killing for a Reason) form a class whose members cannot be identical to any of the members of the class of events (like those in Phobia Killing) whose origins are completely cut off from the elements of your psychology over which you have control. If this much is correct, then it is not the case that an action can be done for reasons and yet the agent ends up performing this very same act for no reason at all.

A final comment. We tend to think that cognitive sophistication is required for a creature's actions to possess deontic properties.<sup>37</sup> It may be the case that you ought (in the normative sense) to wave at your neighbor. But it is not the case that your dog ought (in the normative sense) to wave at your neighbor. And this is not because your dog ought to refrain from waving. It is because talk of ought is out of place. Your dog's behavior might be good but it cannot—because it cannot be done for reasons—be up for normative assessment. Judgment-Sensitive Attitudes Are Necessary explains why this is so. Actions that possess deontic properties demand cognitive sophistication because these actions causally require judgment-sensitive attitudes and having these attitudes demands cognitive sophistication.

## II. CAUSAL DEONTIC PRIMACY

In the previous section, I argued that judgment-sensitive attitudes are causally necessary for actions that ought to be done. In this section, I aim to defend

**CAUSAL DEONTIC PRIMACY:** For any agent, *S*, action, *A*, and judgment-sensitive attitudes,  $J_1 \dots J_n$ , if (i) *S*'s *A*-ing cannot obtain without being caused by *S*'s  $J_1 \dots J_n$  obtaining, (ii) it ought to be that *S*'s *A*-ing obtains, and (iii) the deontic properties possessed by *S*'s  $J_1 \dots J_n$  are explained neither by the deontic properties of *S*'s *A*-ing nor some shared third

<sup>37</sup> For discussion, see Christine M. Korsgaard, *The Sources of Normativity* (Cambridge, UK: Cambridge University Press, 1996), pp. 92–93; and Donald H. Regan, "How to Be a Moorean," *Ethics*, cxiii, 3 (2003): 651–77, at pp. 653–54.

factor, then the deontic properties possessed by *S*'s *A*-ing are explained by the deontic properties possessed by *S*'s  $J_1 \dots J_n$ .<sup>38</sup>

Causal Deontic Primacy is cumbersome. So here is a simplification that I hope does not court confusion: If the deontic properties possessed by a necessary causal antecedent are explained by neither the deontic properties of the causal consequent nor a third factor, then the deontic properties possessed by the consequent are explained by the deontic properties possessed by the antecedent.

Examples may help clarify the idea. Suppose that the deontic properties possessed by a causal antecedent are explained by the deontic properties possessed by the consequent. For example, we might hold on to the Deontic Primacy of Actions, and then claim that, insofar as attitudes have deontic properties, these properties are explained by their instrumental role in the production of action.<sup>39</sup> On this view, if you ought to have a certain belief-desire pair, this is explained by the fact that you ought to perform a certain action and this belief-desire pair is the necessary causal antecedent of this action. Here, Causal Deontic Primacy does not apply. Next suppose that some third factor gave a united explanation of both the deontic properties of the antecedent and the consequent. For example, suppose that if you ought to perform a certain action, this is explained by the fact that it would bring about the uniquely best world available to you.<sup>40</sup> Further suppose that you ought to hold the belief-desire pair that is the neces-

<sup>38</sup>No Control No Rationalization and the notion of judgment-sensitive attitudes specified in the previous section put us in a position to hold that the attitudes causally required for you to act possess deontic properties. Recall, I am assuming that whatever the sense of control we can have over judgment-sensitive attitudes such that they can possess deontic properties is the same sense of control that we need to have over our psychological states such that they rationalize the actions they cause.

<sup>39</sup>Some consequentialists claim that you ought to have certain attitudes and this is explained by the fact that having these attitudes play an instrumental role in causing you to do what you ought. For a recent example, see Christopher Woodard, *Taking Utilitarianism Seriously* (Oxford: Oxford University Press, 2019), chapter 9. For a related view focusing on intentions, see Ulrike Heuer, "Reasons to Intend," in Daniel Star, ed., *The Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press, 2018), pp. 865–90.

<sup>40</sup>Global Consequentialism holds that, for any *X*, it ought to be that *X* obtains if and only if and because the obtaining of *X* would bring about the best outcome. So, for attitudes, actions, or anything else, their deontic properties are directly explained in terms of the value of the worlds where they obtain. For discussion, see Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), pp. 24–28; Shelly Kagan, "Evaluative Focal Points," in Brad Hooker, Elinor Mason, and Dale E. Miller, eds., *Morality, Rules and Consequences: A Critical Reader* (Edinburgh: Edinburgh University Press, 2000), pp. 134–55; and Philip Pettit and Michael Smith, "Global Consequentialism," in Brad Hooker, Elinor Mason, and Dale E. Miller, eds., *Morality, Rules and Consequences: A Critical Reader* (Edinburgh: Edinburgh University Press, 2000), pp. 121–33.

sary causal antecedent of this action, but this is not explained by this pair's connection to an action that you ought to perform. Rather it is explained by the fact that holding this belief-desire pair would bring about the best world available to you. Here too, Causal Deontic Primacy does not apply. Finally, suppose that we ruled out these options. The deontic properties of the causal antecedent are explained neither by the deontic properties of the consequent nor some third factor. For example, suppose that we held that you ought to hold some belief-desire pair which is the necessary causal antecedent of a certain action which you ought to perform, and the fact that you ought to hold this belief-desire pair is explained by facts that serve as evidence for and against the accuracy of these attitudes. Since your evidence for and against the accuracy of actions cannot explain the deontic properties of an action—because actions do not represent things as being a certain way—Causal Deontic Primacy applies. The deontic properties of the attitudes explain the deontic properties of the action.

We can defend Causal Deontic Primacy with two weaker claims. The first of these claims is

*Causal Deontic Inheritance:* For any agent,  $S$ , action,  $A$ , and judgment-sensitive attitudes,  $J_1 \dots J_n$ , if  $S$ 's  $A$ -ing could not obtain without being caused by  $S$ 's  $J_1 \dots J_n$  obtaining and it ought to be that  $S$ 's  $A$ -ing obtains, then it ought to be that  $S$ 's  $J_1 \dots J_n$  obtains.

Causal Deontic Inheritance is a specific version of what is known as deontic inheritance.<sup>41</sup> Put generally, if an agent cannot make one possessor of deontic properties obtain without another—the one performatively entails the other—and the one ought to obtain, then the other ought to obtain. A bit more precisely, for any two possessors of deontic properties,  $P_1$  and  $P_2$ , when  $P_2$  obtaining entails  $P_1$  obtaining, if  $P_2$  ought to obtain, then  $P_1$  ought to obtain. Being permitted and being required are closed under performance entailment.<sup>42</sup>

<sup>41</sup> For recent discussion and defense of deontic inheritance, see Portmore, *Opting for the Best*, *op. cit.*, section 4; and Daniel Muñoz and Jack Spencer, "Knowledge of Objective 'Oughts': Monotonicity and the New Miners Puzzle," *Philosophy and Phenomenological Research*, ciii, 1 (2021): 77–91.

<sup>42</sup> I should mention that there are purported counterexamples to deontic inheritance. See, for example, Alf Ross, "Imperatives and Logic," *Philosophy of Science*, xi, 1 (1944): 30–46; and Frank Jackson and Robert Pargetter, "Oughts, Options, and Actualism," *The Philosophical Review*, xcv, 2 (1986): 233–55. But there are also powerful responses to these purported counterexamples to be found in the references mentioned in the previous note. For an argument that would help to reject actualist objections to deontic inheritance, see Benjamin Kiesewetter, "Instrumental Normativity: In Defense of the Transmission Principle," *Ethics*, cxxv, 4 (2015): 921–46. For discussion of this argument, see the response by Stephen J. White, "Transmission Failures," *Ethics*, cxxvii,



Causal Deontic Inheritance operates at the extensional level. It makes no claim concerning explanation. It claims only that we are licensed to make a certain inference about the possessors of deontic properties when one cannot obtain without the other. However, if causal deontic inheritance applies to two possessors of deontic properties and the deontic properties possessed by the first are explained independently, then the deontic properties possessed by the second must be explained by the deontic properties possessed by the first. This follows from the second claim in our defense of Causal Deontic Primacy, namely,

*Inheritance Calls for Unification:* For any agent,  $S$ , action,  $A$ , and judgment-sensitive attitudes,  $J_1 \dots J_n$ , if causal deontic inheritance applies to  $S$ 's  $A$ -ing and  $S$ 's  $J_1 \dots J_n$ , then either the deontic properties possessed by  $S$ 's  $A$ -ing explain the deontic properties possessed by  $S$ 's  $J_1 \dots J_n$ , the deontic properties possessed by  $S$ 's  $J_1 \dots J_n$  explain the deontic properties possessed by  $S$ 's  $A$ -ing, or some shared third factor explains both.

Inheritance rules out the possibility of a non-overlapping explanation.

Why accept Inheritance Calls for Unification? Because, without a unified explanation, deontic conflicts would be easy to generate. All we would need to do is adjust the non-overlapping part of the explanation while leaving in place the conditions that make Causal Deontic Inheritance apply. By way of illustration, consider

Conflicted: Your apologizing to your friend cannot obtain without being caused by your believing that you wronged your friend. You ought to apologize if and only if and because it would make things go uniquely best. Your apologizing would make things go uniquely best. You ought to believe that you wronged your friend if and only if and because you have sufficient evidence that you wronged your friend. It is not the case that you have sufficient evidence that you wronged your friend.

If we accept Causal Deontic Inheritance, we must deny the possibility of Conflicted. For the stipulations of the case lead directly to deontic conflict. Since it will make things go uniquely best, you ought to apologize. Causal Deontic Inheritance thus tells us you ought to believe that you wronged your friend. But, since you do not have suffi-

---

3 (2017): 719–32; and the reply by Benjamin Kiesewetter, “Contrary-to-Duty Scenarios, Deontic Dilemmas, and Transmission Principles,” *Ethics*, cxxix, 1 (2018): 98–115. It should also be noted that Causal Deontic Inheritance is a logically weaker thesis than deontic inheritance, which makes it more plausible because less vulnerable to these objections. For example, because of the causal connections presupposed, it does not seem to give rise to Ross's paradox. I thank an anonymous reviewer for helping me recognize this last point.

cient evidence that you wronged your friend, it is not the case that you ought to believe that you wronged your friend. Causal Deontic Inheritance and the non-overlapping explanation for the deontic properties stipulated in Conflicted result in you being required to believe and your being required to refrain from believing. To avoid this result, we should accept Inheritance Calls for Unification.

Causal Deontic Inheritance and Inheritance Calls for Unification together secure Causal Deontic Primacy. If Causal Deontic Inheritance applies to a given action and set of attitudes and the deontic properties of these attitudes are explained independently, then Inheritance Calls for Unification tells us that the facts concerning the deontic properties possessed by your so acting are explained by the deontic properties possessed by these attitudes. If the deontic properties of your acting were not explained in this way, we would get a non-overlapping explanation, inviting deontic conflicts.

### III. THE DEONTIC INDEPENDENCE OF ATTITUDES

Causal Deontic Primacy puts us within reach of the idea that what an agent ought to do is explained by the attitudes she ought to have. What is still needed is the claim that the deontic properties possessed by judgment-sensitive attitudes are explained neither by the deontic properties of an agent's actions nor some shared third factor. In this section, I will thus defend the

DEONTIC INDEPENDENCE OF ATTITUDES: For any agent, *S*, action, *A*, and judgment-sensitive attitudes,  $J_1 \dots J_n$ , it is always the case that, when *S*'s *A*-ing cannot obtain without being caused by *S*'s  $J_1 \dots J_n$  obtaining and it ought to be that *S*'s *A*-ing obtains, the deontic properties possessed by *S*'s  $J_1 \dots J_n$  are explained neither by the deontic properties of *S*'s *A*-ing nor some shared third factor.

This claim holds that the deontic properties possessed by an agent's judgment-sensitive attitudes are explained independently—they are explained neither in terms of the deontic properties of the acts they cause nor some shared third factor.

If the arguments of the foregoing sections are sound, we are well on the way to securing the Deontic Independence of Attitudes. Judgment-Sensitive Attitudes Are Necessary tells us that, for any action you ought to perform, a certain set of judgment-sensitive attitudes are causally necessary. Causal Deontic Inheritance then tells us that we ought to have these attitudes. Finally, Inheritance Calls for Unification tells us that the deontic properties of this act and the deontic properties of these attitudes must not have a non-overlapping explanation. We are thus entitled to

*All from Any:* For any agent, *S*, action, *A*, and judgment-sensitive attitudes,  $J_1 \dots J_n$ , if (i) *S*'s *A*-ing cannot obtain without being caused by *S*'s  $J_1 \dots J_n$  obtaining, (ii) it ought to be that *S*'s *A*-ing obtains, and (iii) the deontic properties possessed by *S*'s having  $J_1 \dots J_n$  are at least partly explained independently, then the deontic properties possessed by *S* having  $J_1 \dots J_n$  are completely explained independently.

Given All from Any, we can establish the Deontic Independence of Attitudes by showing that at least part of the explanation for why judgment-sensitive attitudes have the deontic properties they do is independent from the deontic properties of actions and independent from some shared third factor. This is what I will attempt in the remainder. I will start with the deontic independence of one judgment-sensitive attitude—belief—and then use this as a blueprint for the rest of the attitudes.

Beliefs aim at truth. They have a mind-to-world direction of fit. They host their contents as being actual. They present their objects under the guise of being the case. These are all different ways of capturing the same platitude: A belief represents its object as being a certain way, and this representation is accurate or inaccurate. This platitude takes us to a second: A belief is the kind of thing for which an agent can have evidence.

These claims—beliefs are the kind of thing that are representationally accurate or inaccurate and so the kind of thing for which an agent's evidence comes to bear—are uncontroversial. Nevertheless it is worth keeping in mind that this makes beliefs unlike most other things. Beliefs can be representationally accurate or inaccurate. But atoms, apples, and actions cannot be representationally accurate or inaccurate. You can have evidence for or against your beliefs. But you cannot have evidence for or against your atoms, apples, and actions. The contrast between belief and action on this front is worth stressing. Unlike attitudes, actions are not representationally accurate or inaccurate because actions do not represent things as being a certain way. You can have evidence for or against a belief, but you cannot have evidence for or against an action.<sup>43</sup>

To secure the deontic independence of belief from action, what is needed is a claim that certain facts which at least partly explain the deontic properties possessed by beliefs are exclusive to belief. And such a claim is readily available and relatively uncontroversial. The deontic status of an agent's belief is at least partly explained by

<sup>43</sup> This paragraph follows Jerry A. Fodor, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, MA: MIT Press, 1987), pp. 10–11.

the agent's evidence for and against the accuracy of that belief. As Kieseewetter notes, "Evidence provides reasons for beliefs—that much is almost universally agreed upon among philosophers. . . . Only very few. . . would deny that at least other things being equal, a person who is in possession of evidence for  $p$  thereby has an epistemic reason for believing  $p$ ."<sup>44</sup> Few would deny this claim because, concerned only with a partial explanation, this claim remains silent on most controversial matters. It is consistent with non-evidential considerations serving as reasons for belief.<sup>45</sup> It is also neutral concerning the nature of evidence and what the possession relation amounts to.<sup>46</sup>

We are making progress toward establishing the deontic independence of beliefs. The deontic properties possessed by an agent's beliefs are partly explained by her evidence for and against their accuracy. She cannot have evidence for or against the accuracy of her actions. Thus evidence cannot stand as a shared third factor, explaining the deontic properties of beliefs (independently from the deontic properties of actions) and explaining the deontic properties of actions (independently from the deontic properties of beliefs).

Still, to trigger All from Any, we need to show that this partial explanation of the deontic properties of beliefs is not in turn explained by the deontic properties of actions. We need to show that your evidence ultimately, not proximally, explains the deontic properties possessed by your beliefs. And, for most beliefs, this result is easily secured. Most beliefs—even those that figure in the causal explanation of action—are about things other than the deontic properties of actions. A belief that a certain act serves as a means to your end, for example, represents this act as serving as a means. Your evidence is evidence for and against the accuracy of this belief. Accordingly, it would be implausible to claim that the ultimate explanation for the deontic status of this belief (assuming its status is partly explained by the evidence) lies with the deontic properties possessed by the action it is about.

Yet it might be thought that for at least some beliefs the deontic properties of actions have a special significance. Suppose, upon seeing the building ablaze and hearing people scream, you come to believe that you ought to phone for help. And further suppose that, in point of fact, you ought to phone for help. Your belief is accurate, and what makes it accurate is that your phoning actually has the prop-

<sup>44</sup> Benjamin Kieseewetter, "Are Epistemic Reasons Normative?," *Noûs*, LVI, 3 (2021): 670–95, at p. 671.

<sup>45</sup> For an overview of this debate, see the essays in Brian Kim and Matthew McGrath, eds., *Pragmatic Encroachment in Epistemology* (New York: Routledge, 2019).

<sup>46</sup> For an overview of this debate, see Kurt Sylvan, "Epistemic Reasons I: Normativity," *Philosophy Compass*, XI, 7 (2016): 364–76, at section 3.

erty of being what you ought to do. Here one may worry that such beliefs—ones that are accurate because of the deontic properties of actions—threaten the case for the deontic independence of beliefs we have been pursuing.

This worry can be assuaged. For the case for the deontic independence of beliefs does not hinge on the claim that the deontic properties possessed by beliefs are partly explained by the accuracy of these beliefs. Rather it hinges on the claim that the deontic properties possessed by beliefs are partly explained by those facts that serve as evidence for and against the accuracy of these beliefs. The evidence that might serve as reasons to believe that you ought to phone the fire department, for example, include facts that indicate, support, or probabilize the accuracy of that belief.<sup>47</sup>

It is worth driving home the point that the deontic status of a belief about the deontic status of an action is not explained by the deontic status of that action. One way to see this vividly is by noting that agents can be deeply mistaken about their available actions.<sup>48</sup> You can have a belief—one that possesses deontic properties—that you ought to perform some action even when the time of its performance occurs after you cease to exist. Since you cannot have a reason for such an action—it lacks a deontic status—we cannot explain the deontic status of your belief in terms of the deontic status of this action. Return to the example of your seeing the building ablaze and hearing people scream. Suppose your evidential situation is the same. As before, you come to believe that you ought to phone for help. Yet, in this version of the case, coming to have this belief is the unexpected end of your existence. You perish the moment it is formed. You never get the chance to phone the fire department. You never have the opportunity to do anything again. Thus, assuming your belief possesses deontic properties, the properties it possesses cannot be explained by the deontic properties of your phoning the fire department, because

<sup>47</sup> I want to remain neutral on the nature of evidence. So I only mean to be gesturing toward something in the vicinity of indicating, supporting, or probabilizing.

<sup>48</sup> Below, I discuss cases where the agent ceases to exist before the time of action. In this case, there is some action that you normally could do, but because of abnormal circumstances you cannot. Yet this is not the only sort of case. Perhaps more bizarre cases make it clearer that actions cannot play this explanatory role. Suppose, for example, the agent is seriously confused about her abilities—she thinks she can literally reach back into the past, say—and yet she has, because of some twisted worldview, evidence for these abilities. In these sorts of cases, there is some action that the agent (perhaps reasonably) believe she ought to perform but it is not even, in principle, a possible action for her—it is not remotely within her power. This discussion, I should note, was inspired by Kieseewetter's inclusion of a "can" clause in his Evidence Principle. See his *The Normativity of Rationality* (Oxford: Oxford University Press, 2017), section 7.8.

your phoning the fire department does not possess deontic properties.<sup>49</sup> What explains the deontic properties possessed by your belief? The tempting answer to give is that at least part of what explains it are the facts that serve as your evidence for and against this belief's accuracy. And if we say this here, then we should also say it about our first variant of this case—where you do not perish and indeed ought to phone the fire department. Thus it seems that the deontic properties of beliefs are, even in cases where the beliefs are about the deontic properties of actions, at least partly explained independently. They are ultimately explained by the agent's evidence.

Before turning to other judgment-sensitive attitudes, we can rehearse the strategy we have been pursuing: Beliefs represent their objects as being a certain way. They can be accurate or inaccurate. An agent can have evidence for or against the accuracy of her beliefs. When a belief possesses deontic properties, the deontic properties possessed are at least partly explained by her evidence for or against the accuracy of the belief. An agent's evidence provides reasons for her beliefs. But this evidence cannot, as a shared third factor, explain the deontic properties of her actions. Nor can facts about the deontic properties of an action provide the ultimate explanation for why evidence at least partly explains the deontic properties of beliefs, not even when the belief is about an action having certain deontic properties. Thus we can, via All from Any, arrive at the deontic independence of beliefs. What I will next argue is that we can run this same strategy for all judgment-sensitive attitudes.

Are other judgment-sensitive attitudes like beliefs in representing their object as being a certain way?<sup>50</sup> To answer this question, we can start with emotions. According to the received view, emotions have

<sup>49</sup> Does this response—that an agent ought to have some attitude because of misleading evidence—conflict with Causal Deontic Inheritance? No. Causal Deontic Inheritance states only a sufficient condition on deontic property possession for judgment-sensitive attitudes. It ties together the deontic status of an attitude and action when both the attitude and action have a deontic status and the attitude is causally necessary for the action. But the principle is silent when this causal connection does not hold or when the action lacks a deontic status. It is thus consistent with Causal Deontic Inheritance for it to be the case that you ought to hold some belief even if you never perform any action again (and hence the belief is not causally connected to any action that possesses deontic properties).

<sup>50</sup> Let me stress that I am not claiming that all judgment-sensitive attitudes just are representations of things being a certain way. Rather I am claiming that they involve a representational element. For example, views that hold that all emotions are part cognitive and part affective are compatible with the coming argument. For a good example of such a view, see C. D. Broad, "Emotion and Sentiment," *The Journal of Aesthetics and Art Criticism*, xiii, 2 (1954): 203–14.

representational components that can be assessed as accurate or inaccurate.<sup>51</sup> Here is Deonna and Teroni:

[A]s a direct consequence of their being directed at particular objects and connected with types of evaluations, emotions are subject to standards of correctness. . . . In this respect, emotions are similar to many cognitive states such as beliefs and perceptual experiences. All these states have conditions of correctness, i.e., they have a content in the light of which it is possible to assess whether they fit the facts or not. . . . The fact that emotions are assessed as correct or incorrect depending on whether or not they fit the facts has prompted philosophers to talk about them as having the mind-to-world direction of fit—they aim, as it were, at representing the world as it is.<sup>52</sup>

When you admire Kant, for example, the object of the admiration is Kant. And your admiration represents Kant being a certain way. Perhaps when you admire Kant, you represent Kant as having certain evaluative properties—for example, being admirable. Perhaps when you admire Kant, you represent Kant as having the properties that make it the case that he is admirable—for example, intelligent, principled, and systematic. In any case, it will not matter for the purposes of the coming argument so long as you indeed represent him in a way that can be accurate or inaccurate. These same remarks hold for other emotions. When you fear the approaching snake, the object of your fear is the snake, and you represent the snake as dangerous or as having the properties that make it the case that it is dangerous. When you pity yourself, the object of your pity is yourself, and you represent yourself as being pitiable or having the properties that make it the case that you are pitiable.

If emotions and beliefs can be accurate or inaccurate, we have covered most judgment-sensitive attitudes. Arguably the only outlier is intentions. And the case for thinking that intentions can be accurate or inaccurate is perhaps even stronger than emotions. For it is commonly thought that there is some intimate connection between intending and representing that you will perform the act that serves as the object of the intention—for example, intentions are beliefs

<sup>51</sup> For the classic statement of this view, see Justin d'Arms and Daniel Jacobson, "The Moralistic Fallacy: On the 'Appropriateness' of Emotions," *Philosophy and Phenomenological Research*, LXI, 1 (2000): 65–90. For a recent overview, see the opening paragraph of Oded Na'aman, "The Rationality of Emotional Change: Toward a Process View," *Noûs*, LV, 2 (2021): 245–69.

<sup>52</sup> Julien A. Deonna and Teroni Fabrice, *The Emotions: A Philosophical Introduction* (London: Routledge, 2012), p. 6.

or some combination of beliefs with another non-belief state.<sup>53</sup> But, even if we want to stay neutral concerning whether an intention is or involves representing what you will do, we can still claim that an intention represents its object as being a certain way. We can claim, as we did for emotions, that intentions represent actions as being choice-worthy, worthy of pursuit, ought to be done, or desirable. Or they represent actions as having the features that make it the case that they are choiceworthy, worthy of pursuit, ought to be done, or desirable.

We thus seem well positioned to hold that all judgment-sensitive attitudes represent their objects as being a certain way such that they can be accurate or inaccurate.<sup>54</sup> And with this claim, the rest of the strategy we ran for the deontic independence of beliefs falls into place. Just as with beliefs, insofar as a type of judgment-sensitive attitude can be representationally accurate or inaccurate, it is the type of thing for which we can have evidence. The relevant evidence for or against a given judgment-sensitive attitude (for example, admiration) just is the evidence for or against the accuracy or inaccuracy of the attitude (for example, that its object is admirable or that its object has the properties that make it admirable). And this evidence at least partly explains the deontic properties possessed by these attitudes. Here is Deonna and Teroni following up their remarks from the previous passage:

[S]tandards of correctness so conceived should be distinguished from epistemological standards by which we assess the justification of emotions. . . . Bernard has good reasons to be elated if he has just heard from a reliable witness that his wife is in much better health. His reasons would be bad were his elation based on a report from a notoriously unreliable witness. In short, our emotions are sometimes justified, and sometimes unjustified. And they can be unjustified even if, by chance perhaps, they meet the standard of correctness just mentioned. That is the case if, although Bernard's wife is really in good health, his elation is based on the testimony of an unreliable informant. In this last respect, emotions resemble beliefs, for which we also often request reasons (that may or may not justify them), and differ from perceptions that can be said to be correct or incorrect but which are not justified by reasons.<sup>55</sup>

<sup>53</sup> For the view that intending just is believing, see Berislav Marušić and John Schwenkler, "Intending Is Believing: A Defense of Strong Cognitivism," *Analytic Philosophy*, LIX, 3 (2018): 309–40. For the idea that intending involves belief and some non-cognitive component, see H. P. Grice, "Intention and Uncertainty," *Proceedings of the British Academy*, LVII (1971): 263–79.

<sup>54</sup> For further defense of this idea, see Gideon Rosen, "The Alethic Conception of Moral Responsibility," in Randolph Clarke, Michael McKenna, and Angela M. Smith, eds., *The Nature of Moral Responsibility* (Oxford: Oxford University Press, 2015), pp. 65–87; and Portmore, *Opting for the Best*, *op. cit.*, pp. 54–61.

<sup>55</sup> Deonna and Teroni, *The Emotions*, *op. cit.*, pp. 6–7.



The same holds for intentions. When your intention possesses deontic properties, the deontic properties possessed are partly explained by the evidence for and against the intention's accuracy (for example, that the action is choiceworthy or that it has the properties that make it choiceworthy).<sup>56</sup>

Notice that, even if we hold the view that intentions represent their objects as having deontic properties, the remarks made above concerning beliefs about the deontic properties of actions apply. What is being claimed is that the facts that serve as evidence for or against the accuracy of the intention at least partly explain the deontic properties of the intention. And this is crucially different from claiming that the accuracy of the intention partly explains the deontic properties of the intention. Again, we can see this point by looking to contrast cases where the agent is in the same evidential situation, but in one case the agent is mistaken about what it is within her power to do. Since one can have a deontic property possessing intention to perform some act even when she cannot possibly perform it (and so this act cannot possess deontic properties), to explain the deontic properties of the intention we must appeal to the agent's evidence. And if we are forced to appeal to the agent's evidence in this case, then we should appeal to the agent's evidence in all cases.

We have thus arrived at

*Partly Explained by Evidence:* For any agent, *S*, and judgment-sensitive attitude, *J*, if *S*'s having *J* possesses deontic properties, then the deontic properties possessed by *S*'s having *J* are partly explained by *S*'s evidence for and against the representational accuracy of *J*.<sup>57</sup>

<sup>56</sup> For further defense of the various ideas found in this paragraph, see Patricia S. Greenspan, "A Case of Mixed Feelings: Ambivalence and the Logic of Emotion," in A. O. Rorty, ed., *Explaining Emotions* (Berkeley: University of California Press, 1980), pp. 223–50; Gibbard, *Wise Choices, Apt Feelings*, *op. cit.*, pp. 36–40; Thomson, *Normativity*, *op. cit.*, p. 131; M. Smith, "Parfit's Mistaken Meta-ethics," *op. cit.*; Nathaniel Sharadin, "Reasons Wrong and Right," *Pacific Philosophical Quarterly*, xcvi, 3 (2016): 371–99; and Hichem Naar, "The Fittingness of Emotions," *Synthese*, cxcix (2021): 13601–19.

<sup>57</sup> Partly Explained by Evidence suggests the view that all reasons for attitudes are constituted by evidence for accuracy. For recent discussion of this view, see Kiesewetter, "Are Epistemic Reasons Normative?," *op. cit.*, section 1.2. And it seems to suggest the view that all the right-kind reasons for attitudes are constituted by evidence for their accuracy. For an overview of the right-/wrong-kind reasons debate, see Jan Gertken and Benjamin Kiesewetter, "The Right and the Wrong Kind of Reasons," *Philosophy Compass*, xii, 5 (2017): 1–14. Admittedly, these are controversial positions. The central point of contention is driven by cases where evidence for the attitude's accuracy does not seem to be a reason to hold the attitude. For example, a reliable friend's testimony that Kant is admirable may serve as part of your evidence, but it does not seem to be a reason for you to admire Kant. I believe that counterexamples like these can be explained away. But it is worth noting that, strictly speaking, Partly Explained by Evidence does

That the deontic properties of all judgment-sensitive attitudes are at least Partly Explained by Evidence undergirds much of our interpersonal reasoning. If the deontic properties possessed by your judgment-sensitive attitudes are partly explained by the evidence for and against the accuracy of the attitude, then I can reason with you by supplying considerations that indicate, support, or probabilize the accuracy or inaccuracy of your attitudes. We both see a snake but you do not fear it. I might say: "Remember the guide told us to avoid all snakes in this area. Look—do you see its stripes? I know that pattern. It means it's venomous. We are far from medical services. It is heading straight for us!" In reasoning with you, I am trying to provide you with evidence for representing the approaching snake as dangerous or as having the properties that make it dangerous. Evidential considerations are my primary means of persuasion. Partly Explained by Evidence accounts for why this is so.

We are close to the Deontic Independence of Attitudes, but a final objection remains. Return to beliefs about the deontic properties of actions and intentions (granting the claim that intentions represent acts as having certain deontic properties).<sup>58</sup> Above, I tried to show that the deontic properties of these attitudes are not explained by the deontic properties of the actions that these attitudes are about. In cases where the agent is mistaken about the actions she can perform, the deontic properties of the action (because the action is in fact unperformable and so does not possess these properties) cannot be what ultimately explain the deontic properties of such beliefs or intentions. Yet it might be objected that, although this shows that the deontic properties of these beliefs and intentions are not explained in terms of the actions they are about, we should not have been looking to the deontic properties of particular actions in the first place. Instead, we should look to general truths about reasons for action—truths about what things in a given circumstance count in favor of a certain response.<sup>59</sup> Why is some fact evidence for the accuracy of an

---

not imply that all an agent's evidence for the representational accuracy of an attitude serves as a reason for that attitude. It says only that, for any judgment-sensitive attitude, some of the reasons for holding it are provided by the agent's evidence. I thank an anonymous reviewer for pressing me to acknowledge that I seem committed to these controversial positions.

<sup>58</sup> Notice that, if intentions represent actions as having certain deontic properties, the deontic primacy of actions can be false and yet our intentions can still be accurate. For representing an action as having a certain deontic status is to come down on an extensional (not explanatory) issue.

<sup>59</sup> It might be helpful to think of these general truths as what T. M. Scanlon calls "pure normative claims." See his *Being Realistic about Reasons* (Oxford: Oxford University Press, 2014), p. 37. I thank an anonymous reviewer for this suggestion.

intention? According to this objection, if the accuracy standard for intentions is a deontic property of the action, then part of the explanation for why some fact is evidence for an intention's accuracy appeals to general truths about reasons for action. The Moorean, for example, might answer that part of what explains an agent's evidence for an intention's accuracy is the general truth that one always has most reason to do what brings about the best state of affairs.<sup>60</sup>

This objection hinges on the claim that, for attitudes that represent an action as having certain deontic properties, an agent's evidence for and against these attitudes is partly explained by general truths about reasons for action. But we should not accept this claim.

First, this claim puts us in an implausibly strong position when faced with normative skepticism. We are not licensed to infer that there are general truths about reasons for action from our evidence for or against attitudes that represent an action as having deontic properties. If we could, then we could show that skepticism concerning practical reason is mistaken simply by showing that we have any evidence of this sort. Normative skeptics are not so easily defeated.

Second, given the immense differences in the evidential situations of agents, no purported general truths about reasons for action seem positioned to be part of the explanation of all the evidence. To see this, keep the Moorean's purported general truth in mind and consider

The *Attributivists*: Attributivists are an isolated group of people whose societal practices, culture, and education fanatically revolve around teachings in line with the views of Peter Geach, Paul Ziff, and Judith Jarvis Thomson. After many generations, the notions *good-simpliciter*, *valuable state of affairs*, and the like are completely inaccessible, indeed inconceivable, to members of the community.

Suppose that an *Attributivist* intends to keep a promise, representing the keeping of this promise as having certain deontic properties. She has, given her upbringing and experience, ample evidence for this intention's accuracy. Let us also assume that this promise, if kept, will bring about the worst state of affairs. Can the Moorean's purported general truth about reasons for action explain her evidence?

The answer appears to be No. The Moorean's purported truth is unintelligible to the *Attributivist*, and thus completely outside her epistemic ken. We were looking for an explanation for why certain facts

<sup>60</sup> I owe this objection and the wording of it—especially the last few sentences of this paragraph—to an anonymous reviewer.

probabilize, indicate, or support the accuracy of a certain representation for an agent. To look to what is inconceivable for her is to look in the wrong place. What is nonsense for an agent cannot be what explains why representing things as being a certain way is sensible for her. Moreover, her evidence suggests what clashes with the Moorean's purported truth. As described, the Attributivist's evidence points toward representing the keeping of the promise—a promise that if kept will bring about the worst—as what she ought to do. How could the Moorean's purported truth explain the evidence that points against it? Reasons to intend to do what will bring about the worst cannot be explained by reasons to act in ways that will bring about the best. In short, given that the Moorean's purported truth cannot be cognized by her and conflicts with what her evidence suggests, we can conclude that it is not part of the explanation of her evidence.

As we have seen, the Moorean's purported truth does not stand in the right relationship to all agents to explain the total body of evidence they have for and against attitudes that represent certain actions as possessing certain deontic properties. And standing in the right relationship cannot be achieved simply by shifting away from Moore. For any purported general truth about reasons for action (compatible with the objection under discussion), it is possible for there to be sensible agents who are not in a position to cognize it and have evidence for claims that conflict with it. Thus, even if we assume that intentions represent actions as possessing certain deontic properties, it is not the case that part of the explanation for why some fact is evidence for an intention's accuracy appeals to general truths about reasons for action. The explanatory story ends with evidence.<sup>61</sup>

The deontic properties possessed by an agent's judgment-sensitive attitudes are partly explained by those facts that serve as evidence for and against the representational accuracy of these attitudes. This explanation is ultimate. Partly Explained by Evidence thus takes us, via All from Any, to the Deontic Independence of Attitudes.

#### IV. CONCLUSION

We can now present the argument against the Deontic Primacy of Actions in one place:

<sup>61</sup> Does this amount to a denial that there are general truths about reasons for action? No. What the upshot of this discussion shows is that their formulation will differ significantly from Moore's. The Moorean's purported truth picks out certain facts and claims that these facts call for the same response for all agents in all circumstances. If the main thesis of this article is correct, general truths about reasons for action will make essential reference to an agent's reasons for her judgment-sensitive attitudes. I will not try to give a formulation of reasons for action in terms of an agent's reasons for her attitudes here because I believe this is not a straightforward matter.

1. For any agent, *S*, and action, *A*, if it ought to be that *S*'s *A*-ing obtains, then *S*'s having certain judgment-sensitive attitudes is causally necessary for *S*'s *A*-ing to obtain.
2. For any agent, *S*, action, *A*, and judgment-sensitive attitudes,  $J_1 \dots J_n$ , if (i) *S*'s *A*-ing cannot obtain without being caused by *S*'s  $J_1 \dots J_n$  obtaining, (ii) it ought to be that *S*'s *A*-ing obtains, and (iii) the deontic properties possessed by *S*'s  $J_1 \dots J_n$  are explained neither by the deontic properties of *S*'s *A*-ing nor some shared third factor, then the deontic properties possessed by *S*'s *A*-ing are explained by the deontic properties possessed by *S*'s  $J_1 \dots J_n$ .
3. It is always the case that, when *S*'s *A*-ing cannot obtain without being caused by *S*'s  $J_1 \dots J_n$  obtaining and it ought to be that *S*'s *A*-ing obtains, the deontic properties possessed by *S*'s  $J_1 \dots J_n$  are explained neither by the deontic properties of *S*'s *A*-ing nor some shared third factor.
4. Hence, for any agent, *S*, and action, *A*, if it ought to be that *S*'s *A*-ing obtains, the deontic properties possessed by *S*'s *A*-ing are explained by the deontic properties possessed by the judgment-sensitive attitudes causally necessary for *S*'s *A*-ing to obtain.

If this argument is sound, the deontic properties of actions are explained by the deontic properties of judgment-sensitive attitudes.

ANDREW T. FORCEHIMES

Nanyang Technological University