

The Context-Undermining of Practical Reasons

Author(s): Garrett Cullity

Source: *Ethics*, Vol. 124, No. 1 (October 2013), pp. 8-34

Published by: The University of Chicago Press

Stable URL: <http://www.jstor.org/stable/10.1086/671390>

Accessed: 13-06-2018 05:49 UTC

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://about.jstor.org/terms>



JSTOR

*The University of Chicago Press* is collaborating with JSTOR to digitize, preserve and extend access to *Ethics*

# The Context-Undermining of Practical Reasons\*

*Garrett Cullity*

---

Can one fact deprive another of the status of a reason for action—a status the second fact would have had, but for the presence of the first? Claims of this kind are often made, but they face substantial obstacles. This article sets out those obstacles but then argues that there are at least three different ways in which this does happen.

---

When we deliberate about an action, we need to assess the comparative weights of the reasons for and against it. But perhaps the considerations that provide us with reasons can interact in other ways, and deliberation should be sensitive to those too. In a variety of different forms, it has been suggested that one reason can deprive another actual or potential reason of the deliberative weight it would otherwise have—“defeating,” “undermining,” “silencing,” “disabling,” “excluding,” “undercutting,” “nullifying,” “invalidating,” “cancelling,” “bracketing,” or “attenuating” it.<sup>1</sup> The aim of this article is to explain and defend a normative relationship, “context-undermining,” which belongs on that list.

\* My earlier work on this topic benefited from the comments of many colleagues after presentations at Adelaide, Auckland, Christchurch, Melbourne, Reading, Wellington, and the Australian National University. I am grateful to all of them, and also to Doug Portmore and the reviewers and editors at *Ethics* for their acute and helpful criticism of later drafts. A Discovery-Projects grant from the Australian Research Council provided financial support.

1. For “defeating,” see Sean McKeever and Michael Ridge, *Principled Ethics: Generalism as a Regulative Ideal* (Oxford: Clarendon, 2006), 27–28; Robert Audi, *The Architecture of Reason* (New York: Oxford University Press, 2001), 20. (Compare also Roderick M. Chisholm’s contrast between the “defeat” and “balancing off” of values in “The Defeat of Good and Evil,” *Proceedings and Addresses of the American Philosophical Association* 42 [1968]: 21–38.) However,

*Ethics* 124 (October 2013): 8–34

© 2013 by The University of Chicago. All rights reserved. 0014-1704/2013/12401-0002\$10.00

Talk of “depriving a consideration of deliberative weight” is slippery, though: it slides over an important distinction. Given two facts *A* and *B*, one might claim that *A* is a reason why *B* should not feature in deliberation, or that *A* is a reason why *B* lacks the status of a reason for action. *A* might play the first role without the second.<sup>2</sup> If one of your dinner guests has a seizure, the cost of phoning an ambulance should not feature in your deliberation about what to do, but arguably it is still a reason against making the call: it is a comparatively trivial reason, and its triviality is a reason why you should ignore it.<sup>3</sup> But perhaps there are cases of a second kind. If you are a university teacher, you have reasons of professional integrity not to let personal loyalties influence the grades you award, or your contribution to decisions about scholarships or jobs. Here, perhaps, we should deny that facts about personal attachments count in

---

“defeat” is more often used in a way that allows that outweighing a reason is a way of defeating it. For “undermining,” Walter Sinnott-Armstrong, “Some Varieties of Particularism,” *Metaphilosophy* 30 (1999): 1–12, at 5; Audi, *The Architecture of Reason*, 40–46. For “silencing,” John McDowell, “The Role of *Eudaimonia* in Aristotle’s Ethics,” in his *Mind, Value, and Reality* (Cambridge, MA: Harvard University Press, 1998), 3–22, at 17, “Virtue and Reason,” in his *Mind, Value, and Reality*, 50–73, at 55–56, “Are Moral Requirements Hypothetical Imperatives?” in his *Mind, Value, and Reality*, 77–94, at 91, and “Incontinence and Practical Wisdom in Aristotle,” in *Identity, Truth and Value: Essays for David Wiggins*, ed. Sabina Lovibond and S. G. Williams (Oxford: Blackwell, 1996), 95–112, at 102; David Wiggins, “Deliberation and Practical Reason,” in *Essays on Aristotle’s Ethics*, ed. Amélie Rorty (Berkeley: University of California Press, 1980), 221–40, at 234–35; Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge, MA: Harvard University Press, 1993), 37; Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), 52. For critical discussion of McDowell, see Jeffrey Seidman, “Two Sides of ‘Silencing,’” *Philosophical Quarterly* 55 (2005): 68–77. For “disabling,” Jonathan Dancy, *Ethics without Principles* (Oxford: Clarendon, 2004), 41. For “excluding,” Joseph Raz, “Reasons for Action, Decisions and Norms,” *Mind* 84 (1975): 481–99, at 482–88, and *Practical Reason and Norms*, 2nd ed. (Oxford: Oxford University Press, 1999), chap. 1 and Postscript. For “undercutting,” Mark Schroeder, “Holism, Weight, and Undercutting,” *Noûs* 45 (2011): 328–44. For “nullifying,” T. R. Girill, “On the Defeasibility of Duties,” *Journal of Value Inquiry* 12 (1978): 197–209, sec. 4. For “invalidating,” Kalle Grill, “Anti-paternalism and Invalidation of Reasons,” *Public Reason* 2 (2010): 3–20. For “cancelling,” Raz, *Practical Reason and Norms*, 27. For “bracketing,” T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), 52. For “attenuating,” Dancy, *Ethics without Principles*, 42.

2. The facts meeting that description overlap extensively with what Joseph Raz calls “exclusionary reasons,” which are reasons not to act for, or not to be guided by, other reasons. But there is a difference, since including a reason in deliberation but then not acting on it would satisfy Raz’s definition. (See Raz, “Reasons for Action, Decisions and Norms,” 482–88, and *Practical Reason and Norms*, chap. 1 and Postscript.) Compare also Kant’s claim: “pure practical reason does not require that one should *renounce* claims to happiness but only that as soon as duty is in question one should *take no account of them*”; *Critique of Practical Reason* (1788), trans. Mary Gregor, in Immanuel Kant, *Practical Philosophy* (Cambridge: Cambridge University Press, 1996), 5:93; italics in original.

3. “Deliberation about what to do,” I take it, is the conscious reflection on one’s reasons for action that issues either in a decision or in a conclusion about what one ought to do. In this article, my claims about deliberative exclusion will be about considerations that ought to be excluded from deliberation of both kinds.

favor of awarding someone a higher grade. In other contexts, such facts give you reasons for action: here, they do not.

If there are cases of the second kind, that seems important. To have an accurate picture of the ways in which our reasons can interact to determine what we should do, we need to know whether these include the second sort of relationship. But if so, we need to answer some obvious questions. Why not say instead that you do have a reason to benefit your relatives but should exclude it from deliberation? Why recognize any cases of the second kind? Or the first, for that matter?<sup>4</sup> Having made the distinction between the two kinds of case, how exactly should we apply it? These questions are embarrassing. The literature often fails to make this distinction at all. When it does so, it fails to offer a criterion for applying it.<sup>5</sup> And it fails to notice just how hard it is to show that there really are any cases of the second kind at all, rather than simply the first. Those are the gaps this article tries to fill.

To make the discussion manageable, I focus on cases of “context-undermining.” Fact *A* bears this relationship to *B* (I stipulate) when (i) *A* deprives *B* of the status of a reason for action and (ii) *A* does not do that by showing that *B* is a fact about an attitude with a misdirected content.<sup>6</sup> The point of condition (ii) is to set aside cases like prospective pleasure, which (it is often claimed) loses its reason-giving status when its content is sadistic.<sup>7</sup> I think that kind of case (“content-undermining”) needs a different kind of explanation.<sup>8</sup> Instead, I concentrate here on cases like the professional one described above. There, your personal at-

4. In characterizing virtue in terms of the “silencing” of reasons, John McDowell suggests we should see a case like the first—where a virtuous person will not consider the cost of helping—as one in which there is no such reason. To be in a state of full Aristotelian virtue, rather than mere continence, he maintains, “is to see the relevant reasons for acting, on occasions when they coexist with considerations that on their own would be reasons for acting otherwise, as, not overriding, but silencing those other considerations—as bringing it about that, in the circumstances, they are not reasons at all” (“The Role of *Eudaimonia* in Aristotle’s Ethics,” 17). This does not entail that they are *not* reasons; but it does entail that a fully virtuous person must think they are not.

5. The main exception to these generalizations is Raz’s prominent work on exclusionary reasons. For his distinction, see nn. 2 and 16; for his criterion, n. 12.

6. Perhaps there can be cases of partial context-undermining, where *B*’s status as a reason is diminished, but not entirely lost. But in this article, for clarity, I concentrate on cases in which (I argue) context-undermining is complete.

7. For this claim, see, e.g., Dancy, *Moral Reasons*, 56; Luke Robinson, “Moral Holism, Moral Generalism, and Moral Dispositionalism,” *Mind* 115 (2006): 331–60, at 337; David McNaughton, *Moral Vision: An Introduction to Ethics* (Oxford: Blackwell, 1988), 193; McKeever and Ridge, *Principled Ethics*, chap. 2; and Frank Jackson, Philip Pettit, and Michael Smith, “Ethical Particularism and Patterns,” in *Moral Particularism*, ed. Brad Hooker and Margaret Olivia Little (Oxford: Clarendon, 2000), 79–99, at 97.

8. Some philosophers think that there is no phenomenon of “content-undermining,” since the fact that you will gain pleasure from an activity is never a reason in its favor: the reason is always more fine-grained than that. See, e.g., Roger Crisp, “Particularizing Partic-

tachments fail to provide you with reasons for certain actions, not because they are themselves attachments to the wrong things, but because the context is not one in which they function as reasons.

Or so I claim. Section I begins by describing the strong-looking obstacles to admitting any cases of context-undermining at all. The rest of the article sets out my response. It is structured around a description and explanation of three main ways in which context-undermining can arise. I do not claim that these are the only ways; but by examining them, we will learn what needs to be done to defend the existence of others.

## I. THE PROBLEM

A pair of simple arguments might seem to make the acceptance of context-undermining the default view, with its opponents carrying the burden of justification. First, we have the analogue for beliefs. Ordinarily, the fact that something appears black to you is a reason for believing that it is black, but if you know you are looking at a tomato under blue light, the appearance-fact loses that status. Since this can be true of reasons for belief, why not of reasons for action?<sup>9</sup> And second, there can clearly be reasons *for* other facts' having the status of practical reasons; so why not reasons against?<sup>10</sup> To adapt an example from Hume: the fact that it keeps you in good health can be a reason for getting regular exercise, because it enables you to work productively.<sup>11</sup> Here, there is an action (exercising), a first-order reason for the action (promoting your health), and a second-order reason (assisting your work) why the first counts in favor of the action. Should we not then expect there to be second-order reasons against as well as for? Decisive reasons of that kind will give us cases of undermining.

However, these arguments conclude with questions, and there are replies. Why not admit the existence of context-undermining? There are three immediate answers. The first concerns explanatory simplicity. Whatever the reasons are for excluding personal attachments from institutional decision-making, they can be recognized while accepting that you have *a* reason to favor your relatives. Usually the former reasons outweigh the latter. So we can explain how you ought to think and act without needing to introduce the apparatus of normative relationships of

---

ularism," in Hooker and Little, *Moral Particularism*, 23–47, at 37; and Joseph Raz, "The Truth in Particularism," in Hooker and Little, *Moral Particularism*, 48–78, at 61.

9. See Dancy, *Ethics without Principles*, 49–51; Scanlon, *What We Owe to Each Other*, 52.

10. See Dancy, *Ethics without Principles*, 41.

11. David Hume, *An Enquiry Concerning the Principles of Morals* (1751), ed. L. A. Selby-Bigge and rev. P. H. Nidditch, 3rd ed. (Oxford: Clarendon, 1975), 293. Whereas Hume introduces his example with an explanatory question—he wants to "ask a man *why he uses exercise*"—I am asking a normative one. (Italics in original.)

undermining in addition to those of outweighing. This approach is consistent with the phenomena and simpler, so it is better. Second, an outweighing explanation appears more consistent with the phenomena than an undermining explanation, in the following way. Suppose you do what you should not—make a nepotistic job appointment, say. That might be wrong but would not be totally irrational: it would be comprehensible in a way that, say, picking someone at random from the phone book and appointing them would not. There would be something good about it, namely, the help you give your relative.<sup>12</sup> So that suggests that there is a reason for the nepotistic action. Moreover (third), suppose we change the example by making the consideration of personal attachment more dramatic. A deranged person has threatened to kill your child unless you appoint the weakest candidate on the job shortlist. Surely that makes a difference to how you should think and act. And the most straightforward way to account for this is to say that your personal attachments always give you reasons for action, with the strength of those reasons depending on the magnitude of the interests at stake.<sup>13</sup>

So the initial arguments for accepting the existence of context-undermining are met by forceful counterarguments: from simplicity, from doing what you should not, and from variation in weights. Moreover, there is a more serious problem: the two initial arguments actually backfire.

The flaw in the first argument is that the explanation of its plausibility as applied to reasons for belief blocks its application to reasons for action. The facts that are evidence for our beliefs provide reasons for them by making them likelier to be true. Your knowledge about the lighting makes a difference to whether the tomato's appearing black is a reason for believing that it is black because it makes a difference to whether its having that appearance makes it likely that it is black. So it is because the relationship of support between evidence and belief is a probabilizing relationship that we can explain the undermining of reasons for belief. But the way normative practical reasons support actions is not by making them likelier to be true.

The second argument exposes a deeper problem. To see this, it helps to introduce some terminology. "Response-reasons," we can say, are reasons for making responses of various kinds to the world: responses of action, belief, feeling, and so on. "Determinative reasons" are reasons

12. Related to this are the grounds Raz gives for thinking that exclusionary reasons do not extinguish the reasons they exclude: it makes sense to regret not acting on an excluded reason, to feel some "pull" toward acting on it, and to think it is a matter of good luck if one does manage to act in conformity with it (*Practical Reason and Norms*, 41–45, 184–86).

13. A parallel objection is made by commentators on Raz's application of exclusionary reasons to the law. See, e.g., D. S. Clarke Jr., "Exclusionary Reasons," *Mind* 86 (1977): 252–55; Frederick Schauer, *Playing by the Rules* (Oxford: Clarendon, 1991), 88–93.

for something's being the case.<sup>14</sup> A reason for deliberative exclusion is a second-order response-reason—a reason for the response of excluding some other reason from deliberation.<sup>15</sup> Second-order response-reasons can come in various other forms. You might have a reason not to regard some fact as a reason, even though it is; not to be guided by it in acting;<sup>16</sup> to believe falsely that it is not a reason; to induce in yourself a false belief of this kind;<sup>17</sup> to avoid thinking about some reason-giving fact; and so on.<sup>18</sup> By contrast, what a context-underminer is a reason for—another fact's losing the status as a reason—is not a kind of response: it is the absence of a normative relationship between a fact and a response. A context-underminer is a second-order reason, since it takes another reason as its object.<sup>19</sup> But it is a second-order determinative reason, not a response-reason.

The exercise example illustrates one kind of second-order determinative reason: a reason *for* some other fact's being a reason. The fact that good health will assist your work is a reason why promoting your health is a reason for getting exercise. But notice that what makes that true is a relationship between two response-reasons. It is because there is a reason for the response of promoting your health, and exercising is a way of doing that, that there is a reason for the response of exercising. The problem is that there is no ready analogue of this explanation for context-underminers: reasons against some other fact's being a reason. When I claim that facts about your institutional role are reasons why

14. "Determinative reasons" is not an ideal label for this class of reasons, because it includes reasons for something's being the case which do not (because they are outweighed) actually succeed in making it the case. But "potentially determinative reasons," although more accurate, is too cumbersome. Compare Dancy, *Ethics without Principles*, 33.

15. Sometimes it seems that there are second-order reasons permitting but not requiring us to exclude first-order practical reasons from deliberation: for examples, see Patricia Greenspan, "Asymmetrical Practical Reasons," in *Experience and Analysis*, ed. Maria E. Reicher and Johan C. Marek (Vienna: OeBV & HPT, 2005), 115–22.

16. Raz, *Practical Reason and Norms*, chap. 1 and Postscript. A Razian exclusionary reason does not alter the balance of reasons for or against an action. In Raz's original (1st ed., 1975) presentation, he describes exclusionary reasons as reasons not to act for other reasons. However, as is pointed out by G. F. Schueler, "Exclusionary Reasons," *Personalist* 60 (1979): 407–10, that seems to make every contrary first-order reason an exclusionary reason. The formulation given in the 1999 Postscript, where exclusionary reasons are described as reasons not to be guided by other reasons, therefore seems preferable.

17. This might happen in a "Schelling case," where the most effective way to prevent someone from threatening you is to convince him you are irrational. See Thomas C. Schelling, *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960).

18. "Deliberative exclusion" as described in n. 3 does not seem equivalent to any of the members of this list. You could regard a fact as a reason but not include it in your deliberation (Raz gives examples of that—see n. 2); and you could include it in deliberation without being guided by it in acting—for example, by thinking it is outweighed.

19. More carefully: it takes a consideration which would otherwise have reason-giving status and removes that status from it.

your personal attachments are not reasons for benefiting your relatives, what makes that true cannot be a relationship between two response-reasons. The claim, after all, is that your personal attachments are not response-reasons in this context. But if that could not make it true, what could?

It is tempting to reply: the facts about your institutional role are reasons for the response of behaving impartially, and you behave impartially by not allowing your personal attachments to enter your deliberation as reasons bearing on your exercise of institutional office. But that does not help. It delivers only reasons for deliberative exclusion, and those are second-order response-reasons, not context-underminers.<sup>20</sup> Our problem is to get beyond the former to the latter.

So, far from finding that the acceptance of context-undermining is the default view, we have uncovered an awkward challenge. The two apparent arguments in its favor backfire, and there are three counter-arguments. Not only does the burden of justification lie with the proponent of context-undermining, but it is hard to see how it can be discharged. Showing why a fact should not be admitted as a reason within one's deliberation is not good enough. And without a compelling argument for accepting context-undermining, we face three arguments for rejecting it.

## II. ASSUMPTIONS

In addressing this challenge, I shall make six main assumptions about normative practical reasons.

- (i) Such reasons are facts which bear a normative relation to types of action. We can refer to that relation as "counting in favor."<sup>21</sup> (The reason is the fact, not the relation; the fact's being a reason is its bearing that relation to a type of action.)<sup>22</sup>
- (ii) We do have such reasons, and we are capable of being wrong about them.
- (iii) Morality gives us such reasons, whether we recognize them or not, and those reasons can be decisive.

20. This seems true of McDowell's defence of "silencing." Virtuous thought and agency may involve excluding some considerations from your deliberation; but that falls short of showing that they are not reasons. See n. 4.

21. Compare, e.g., Raz, *Practical Reason and Norms*, 11–12; Scanlon, *What We Owe to Each Other*, 17; Dancy, *Ethics without Principles*, 15, 29. According to Joshua Gert, "Normative Strength and the Balance of Reasons," *Philosophical Review* 116 (2007): 533–62, there are two different counting in favor relations—counting in favor of permission and counting in favor of requirement. I remain agnostic about that here.

22. Often, the actions for which there are reasons do not happen. So the actions for which those reasons are reasons must be act-types, rather than particular events. However, that is consistent with thinking that the relevant act-type is highly specific.

- (iv) Our reasons collectively require some actions and permit others. When I say that a reason is “decisive” I mean that it makes the first of those things the case, and when “sufficient” the second. The notions of being required or permitted by our reasons I leave unanalyzed.
- (v) The facts that supply us with the practical reasons about which we can reach verdicts of those two kinds are determinative reasons for the truth of those verdicts.<sup>23</sup>
- (vi) “Ought” and “should” can be used to express verdicts about what our reasons for and against an action require, all things considered. I use them in that way here. (Given (iv)–(vi), a decisive reason for action is a reason why you ought to do it.)

These assumptions are widely made but sometimes contested. Their importance for the subsequent argument varies. As far as I can see, little of this discussion would survive the rejection of (ii) or (iii)—the topic itself would then be of little interest—but it could be reworded to accommodate alternatives to the other assumptions.<sup>24</sup>

Since writers who make assumptions (i)–(vi) often accompany them with other, more contentious ones, I add some disclaimers. Assumption (i) provokes questions about the conceptual reducibility and the metaphysics of the “counting in favor” relation, and about the nature of the facts and responses that stand in that relation. I take no sides on those questions here. In (iii), I am not assuming that the reasons morality gives us are always decisive. I accept that “ought” and “should” have important uses apart from the one identified in (vi). I remain agnostic about the conceptual priority-relation between reasons and the verdictive “ought.”<sup>25</sup> And, importantly, I make no assumption about the way in which the weights of reasons contribute to the two kinds of verdict mentioned in (iv). I do not assume that there is a single dimension along which we can measure “the” strength of our reasons; nor that what we ought to do is what there is “most” reason to do.<sup>26</sup> I am not ruling out Joseph Raz’s view, that there are exclusionary reasons not to be guided by certain first-

23. My use of “verdict” in this context follows Philippa Foot, “Are Moral Considerations Overriding?” in her *Virtues and Vices* (Oxford: Blackwell, 1978), 181–88, at 182.

24. The rewording would be too cumbersome, though: hence the assumptions. Thus, to avoid (i), we would need to replace “counts in favor of” throughout with “bears a positive normative reason relation to,” “fact” with “reason-giving consideration,” and so on.

25. For the attempt to analyze normative practical reasons in terms of “ought,” see John Broome, “Reasons,” in *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, ed. R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith (Oxford: Clarendon, 2004), 28–55, at 29–42.

26. For criticism of those claims, see Gert, “Normative Strength and the Balance of Reasons,” esp. sec. 4. Assumption (iv) is a datum for which various candidate explanations offer themselves: here, I remain neutral between them.

order reasons for and against action, that (undefeated) exclusionary reasons defeat the first-order reasons they exclude, and that one ought to act for an undefeated reason.<sup>27</sup> I do not assume that if you have only two options, and you have a reason favoring the first and no reason favoring the second, then you ought to choose the first. (I leave open that your reasons might still permit doing either.) And I make no assumptions about the kinds of incommensurability our reasons can or cannot have.

### III. THREE SOURCES OF CONTEXT-UNDERMINING

My defense of context-undermining will focus on relationships of three kinds: those that arise through occupying a role, through making a commitment, and through the “meaning” of a decision. Each is now introduced with an example accompanied by a description of how, as I see it, the relationship of context-undermining arises. The rest of the article then asks what would be wrong with reading the examples differently.

1. *Role-based* context-undermining. This is illustrated by the opening example, in which an institutional role restricts the reasons that can bear on how you discharge that role. Here, in acting as its role bearer, you exercise agency on behalf of the institution, acting on its authority. It acts through you. Since you are performing its actions, the reasons bearing on what you do are those that are relevant to *its* agency, and those are determined by its purposes and priorities, not your own personal ones. The standards for the exercise of institutional agency are fixed by it, not you—either through explicit rules or implicitly, by reference to the nature and purposes of the institution and the subsidiary function of the role in furthering the institution’s aims.

On this way of looking at it, the role bearer’s relationship to the institution is fiduciary. It is like the relationship a trustee bears toward the beneficiaries of a trust, in being committed to acting for their benefit, not her own. When, in acting as their agent, she invests their money, it is their interests, not hers, that provide the reasons determining what she should do.

Having spoken of acting on the “authority” of the institution, and of the institution as “fixing” or “determining” the role-bearer’s reasons, I should add some qualifying remarks. Few institutions have authority

27. See Raz, “Reasons for Action, Decisions and Norms,” and *Practical Reason and Norms*. However, Raz also emphasizes that “the fact that [first-order reasons] are excluded by an exclusionary reason merely means that they should not be complied with, not that they should not be conformed to”—where conformity with a reason is doing what it is a reason to do, and compliance is being guided by that reason to do it (*Practical Reason and Norms*, 185). So what he means in saying that one ought to act for an undefeated reason is not that one ought to do what is supported by the balance of nonexcluded reasons; it is that one’s action ought to be guided by an undefeated reason.

over whether you participate in them, or authority to command arbitrary obedience from you if you do. The authority institutions have over their role bearers is the authority of authorship: the authority, through their rules and purposes, to settle what counts as acting on their behalf.<sup>28</sup> This falls short of determining what, all things considered, you ought to do. It cannot be that, since you can commit yourself (unwisely) to roles in two institutions that make conflicting demands—such conflicts presuppose that both institutions have the authority to make those demands. But although it falls short of determining what you ought to do, it is more than just the addition of a further reason to those you would otherwise have had. An institution's authority over its role bearers is its power to determine what counts as acting on its behalf—and this includes determining the reasons that bear on the actions that are performed on its behalf.

Saying this leaves many questions open. In particular, it leaves open, for any given institution, (i) whether you should be acting as its agent, (ii) whether it should be subverted or infiltrated for purposes contrary to its own, and (iii) what its reasons really are.<sup>29</sup> We might ask: What settles the answers to (i)–(iii), and what is the relationship between those answers? (Are there cases in which an institution should not be doing something, but nonetheless I should be acting as its agent?) But pursuing those questions here would be a digression: many answers are compatible with the claims I have just made, which are the important ones here. I assume only that sometimes, the answer to (i) is Yes, the answer to (ii) is No, and we can know the answer to (iii)—and that usually, your situation in making institutional appointments falls within that range of cases.

2. *Commitment-based* context-undermining. For a second example, suppose you are to chair a hearing of the Truth and Reconciliation Commission at which I shall be testifying as a witness. I approach you beforehand and ask you to promise not to take my feelings into account in your conduct of the hearing, which you do. As the hearing progresses, you can see that I am becoming distressed. Should you keep your promise? If so, it functions as a context-underminer. In its absence, my distress would have given you a reason to adjourn the hearing; now it does not.

Here again, it is natural to use the language of authority in explaining how the promise affects the normative structure of the situa-

28. Contrast the more substantial kinds of authority discussed in Joseph Raz, ed., *Authority* (New York: New York University Press, 1990).

29. Do institutions really have reasons? If you doubt that, you will need to translate the phrase “the institution's reasons” as “the reasons an institutional role-bearer has insofar as she should be furthering the aims and purposes of the institution.” The substance of the argument would be unaffected by that.

tion.<sup>30</sup> But this time, the context-undermining is not role-based but commitment-based: the restriction on your reasons is imposed not by an authority in whose name you are occupying an institutional role, but by an authority you have given me through an act of voluntary commitment. A valid promise confers a (circumscribed) power to command the performance of what is promised. This is not to say that promises must always be complied with, no matter what. Rather, the claim is that a valid promise grants a form of authority to the promisee; whether that (or any other) authority should be complied with depends on all of the reasons for and against doing so. But if your promise should be complied with, and its content concerns what you will take into account in acting, what you should be complying with is my authority over that.<sup>31</sup> In the example, when you keep your promise you recognize a normative power I have: a power to determine whether facts about my welfare, which have the status of reasons for me, are to have that status for you.<sup>32</sup> To have this normative power is to call for a certain kind of respect. It is to be the owner of my own welfare, having authority over its availability as a reason for others.

In a way, this example gives us the reverse of the kind of fiduciary relationship involved in role-based context-undermining. If you agree to act as my trustee, then my interests replace yours as reasons for acting. In contrast, if you make a promise of the kind just described, the opposite happens: my interests are then removed from the set of your reasons for acting.<sup>33</sup>

As before, I acknowledge but pass over the deeper questions that this invites about the grounding and extent of this authority. An attractive view traces the authority conferred on promisees to another, more fundamental kind of authority: an authority that each of us has to

30. For discussion of the way in which the practice of promising serves our “authority interests,” see David Owens, “A Simple Theory of Promising,” *Philosophical Review* 115 (2006): 51–77.

31. This example is similar to Raz’s example of Colin, who promises his wife to consider only his son’s interests when deciding which school to send him to—thereby excluding reasons such as the financial hardship that other families will undergo if they follow his lead (“Reasons for Action, Decisions and Norms,” 486–87, and *Practical Reason and Norms*, 39). But notice the following difference. In my example, the promisee (I claim) possesses the authority to determine whether his own interests are reasons that bear on the promisor’s action. It is not so clear that Colin’s wife possesses the authority to determine whether someone else’s interests are reasons that bear on Colin’s action. This distinction will be important later.

32. On the normative powers associated with promises, see the exchange between Neil MacCormick and Joseph Raz, “Voluntary Obligations and Normative Powers,” *Supplement to the Proceedings of the Aristotelian Society* 46 (1972): 59–102.

33. I am not making the dubious claim that promises generally function as context-underminers: the explanation I give of context-undermining in relation to promises of this special kind does not extend to others.

require respectful treatment from others, which includes among many other things fulfilling their undertakings to us.<sup>34</sup> This deeper authority is not given to us by anyone: I have it prior to any interactions with you, and it is a background to the possibility of any binding commitments between us. Different ways of filling that out are consistent with the claim I advance here.<sup>35</sup>

3. *Meaning-based* context-undermining. The third example is this: you must decide which one of your two sons to take on a train trip. Suppose Bill loves trains, Ben hates them, and it is Bill's birthday; those facts could count decisively in favor of taking Bill. However, adding a further fact could change the normative structure of the situation significantly. The following, horrible fact would do so: whoever is left behind will die. Now there is no longer a decisive reason to take Bill in preference to Ben. We might insist that there is still a sufficient reason to take Bill. But the only plausible candidate for that reason is that it will save his life. The facts that originally counted in favor of taking Bill—that he loves trains, and it is his birthday—lose that status in this new context. Their status as reasons for taking Bill and leaving Ben is undermined by the fact that this is now a matter of life and death.<sup>36</sup>

The explanation of this case does not follow either of the first two patterns. Rather, we should say this. In one way, adding the further fact does not change the decision you face: you still have to decide which son to take on the train. But obviously, it does change it in another, highly significant way: you are now deciding which of your children's lives to save. Given that your decision has acquired this significance, or "meaning," it loses others.<sup>37</sup> In particular, it is no longer a decision concerning a recreational outing.<sup>38</sup> To a decision with *that* meaning, considerations of Bill's enjoyment and his birthday are relevant as reasons. But with the change of meaning, that ceases to be true. Now only one question calls to be answered: How can you properly respect the value of your

34. One might draw from that view the conclusion that I have the authority to require you not to take my interests into account, whether you have promised to or not. The discussion here is compatible with that stronger claim—but also with its denial.

35. For one noteworthy recent proposal, see Stephen L. Darwall, *The Second-Person Standpoint: Morality, Respect, and Accountability* (Cambridge, MA: Harvard University Press, 2006); for his discussion of promising, 203–9.

36. Compare the discussion of "irrelevant utilities" in F. M. Kamm, *Morality, Mortality*, vol. 1, *Death and Whom to Save from It* (New York: Oxford University Press, 1993), 146.

37. For this way of talking of the "meaning" of a decision or practice—but not the claim about context-undermining—compare Michael Walzer, *Spheres of Justice: A Defence of Pluralism and Equality* (New York: Basic Books, 1983), chap. 1.

38. More carefully: there is an objective sense of the "meaning" of the decision on which it is no longer a decision about a recreational outing if it is a decision about life and death. Were you (and Bill) to persist in thinking that your train trip was a recreational outing, you would be mistaken.

sons' lives? What makes this question so difficult is that, on the one hand, the death of both children is obviously worse than the death of only one; but on the other hand, choosing to save one child's life above the other is bad too. What is the right answer? I don't know. But what does seem plain is that considerations about Bill's enjoyment and his birthday make no contribution to supporting that answer, whatever it is. It would be grotesquely absurd to cite those facts in explaining to someone why you saved Bill rather than Ben. To do that would betray a lack of appreciation of the significance of your decision. So in this third case, context-undermining is meaning-based: the meaning of your decision constrains the reasons that can properly bear on it.<sup>39</sup>

How is context-undermining possible? We have three proposed answers. The common feature, I suggest, is this. Changes in contextual facts can change the nature of an action, and the reasons for an action depend on its nature. We are familiar with this phenomenon in its positive form: contextual facts can make taking your hat off a gesture of respect and, thus, create a reason to do it. In the three cases just discussed, we have negative counterparts of this phenomenon. In the first case, the contextual change alters the agency being exercised; in the third, it alters the meaning of what is done. In the second, there is a change in the proper boundaries of the exercise of agency: respecting me requires recognizing my control over those boundaries, allowing me to reserve my welfare within the sphere of my agency rather than yours. When the nature of an action changes in one of these ways, a fact can lose its normative relevance to the action, and thus its status as a reason.

#### IV. PERSISTENCE

This gives us three proposed answers to the question, "Why does *A* undermine *B*'s status as a reason?" But how convincing are they? In each case, a straightforward challenge needs to be met. Why not say instead that *B* is a persisting practical reason, and *A* is a reason why *B* should be excluded from deliberation? We can call that a Persistence View of the structure of the cases.

A Persistence View of the first case can agree that an institution can have authority over how you conduct yourself in discharging an office-

39. How does this differ from the other cases? Might we not say that taking on an institutional role or making a promise affects the subsequent "meaning" of your decisions too? We could say that. But context-undermining is meaning-based when it is explained by the way in which, through acquiring one meaning, a decision loses others, and that makes certain facts irrelevant. In the other cases, the explanation travels in the opposite direction. It is because personal attachments are irrelevant to the exercise of institutional agency that using an institutional role to benefit one's relatives has a nepotistic meaning.

bearing role within it. When you are making an appointment decision, the institution instructs you (implicitly or explicitly) to appoint the best-qualified candidate and to exclude from consideration factors other than the candidates' qualifications. But on a Persistence View, you do still have *a* reason to benefit your less-qualified relatives. That is a reason against complying with the institutional instruction—a reason that is normally outweighed. The institutional reason is normally decisive, so you should exclude the personal reason from your deliberation. But there is no need to say that the personal reason disappears.

In the second case, you promise not to take my feelings into account. So if your promise ought to be kept, it gives you a decisive reason to exclude my distress from your deliberation about whether to adjourn the hearing. And if so, you have a decisive reason not to perform certain actions: namely, actions of adjourning-the-hearing-in-order-to-spare-me-distress. But a Persistence View can accept this while holding that my distress remains *a* reason to break the promise—an outweighed reason. Although you can give me authority over what you admit into your deliberation, that authority does not extend to being able to make your practical reasons themselves disappear.

A Persistence View of the third case can also agree that the life-and-death “meaning” of the decision about the train trip would make it absurd to resolve it on the grounds of Bill’s birthday enjoyment. But that is a point about the role this reason should play in decision making, not whether it is a reason at all. It is a reason, but not a good enough reason for choosing Bill over Ben, when their lives are at stake. It is like the cost of calling the ambulance in the dinner guest example. The point about the meaning of the action, it can be admitted, has explanatory force—but what it explains is why it would be frivolous to make the decision for this reason, not why it ceases to be a reason at all.

There are other, more radical ways of resisting my treatment of the three cases. It might be denied that there should be any deliberative exclusion in these cases, that promises and institutional roles give us reasons at all, and so on. But here our interest is in addressing the opening problem—the problem of seeing how to provide a compelling case not just for deliberative exclusion but context-undermining. So the opposition to consider comes from the Persistence View. It accepts that these are plausible cases in which a consideration ought to be excluded from deliberation but shows how to account for this without context-undermining. Instead, it offers a “dual-verdict explanation.” A persisting reason is outweighed, and this explains two different verdicts: a verdict about what you should do (appoint the best-qualified candidate, not adjourn to spare me distress, not take Bill for the sake of his enjoyment) and a verdict about how you should deliberate. What is special about these cases is that if you should, all things considered, obey the institution, keep

your promise or be guided by the “meaning” of the decision, then doing those things involves excluding some of your reasons from deliberation. The verdict about action implies a verdict about deliberation. But that does not prevent the excluded considerations from being among the reasons that determine those two verdicts. This approach is theoretically simpler and has the other apparent advantages we noticed in Section I, so it should be preferred.

## V. UNDERMINING VERSUS PERSISTENCE

The Persistence View seems to explain our examples more simply than the interpretation I proposed in Section III (“my interpretation,” as I’ll refer to it from now on). But this advantage disappears, I shall now argue, when we consider a broader range of variations on the examples. We need to ask which interpretation best explains the pattern of normative verdicts that we find across that range—verdicts concerning both action and deliberation. The Persistence View ties the two together; as such, its dual-verdict explanation is tightly integrated and simple. I shall argue that that turns out to be an explanatory liability, not an advantage. It ties different kinds of normative verdicts too tightly together to be plausible. To see this, we need to work through the three cases once more. Each time, I shall present a set of variations on the earlier examples and then explain why the Persistence View’s treatment of them is unsatisfactory.

1. We can start with a pair of variations on the role-based case.

- (a) Your family faces a financial crisis unless you acquire \$50,000 very soon. You have set aside Monday to assess a batch of scholarship applications. According to the scholarship allocation system, making your choice by the end of Monday will secure the best candidate; if you delay a day longer, you will only get the tenth best. On Monday morning, you receive an unexpected opportunity to spend the day acquiring \$50,000. You do so; the tenth best candidate is appointed.
- (b) Your family faces a financial crisis unless you acquire \$50,000 very soon. You look for an opportunity to make \$50,000, but none arises. Although you study the applications conscientiously over several days, your preoccupation makes it difficult to concentrate, and the candidate you appoint, Smith, is only the tenth best. However, you are surprised to find that Smith’s appointment turns out to be the condition for triggering the distribution to you of a \$50,000 legacy by a distant relative.

There is a similarity between these cases: in both, it can be said that helping your family depended on not appointing the best candidate. But there is also a difference, and it seems significant.

One way to bring out its significance is by asking these two questions:

- (1) In (*a*), should you pursue the business opportunity?
- (2) In (*b*), should you appoint Smith?

The answer to (1) could be Yes. Here, your responsibilities to your family and your institution pull in opposite directions; the first could be stronger. If your family's plight is bad enough, the reason to help them could be decisive, outweighing the reason of institutional responsibility which counts the other way. But the answer to (2) is No. The job you have done there in assessing the applications is flawed: you should have appointed the strongest candidate. That is not to say that you are blameworthy for this failure, since it was not due to any lack of diligence. But the "should" of assumptions (i)–(vi) is not the "should" of blameworthiness: it is the "should" of verdicts about decisive objective reasons—the reasons we have whether we recognize them or not. That is the sense in which you should have stayed at home on the day there was an unexpected power outage at work and your day was wasted. It is in that sense that, as the scholarship officer, you should have appointed the strongest candidate, not Smith. You have made a mistake.

So we have two questions, and two answers to them. Now we should ask how successful the two rival interpretations are in explaining those answers.

My interpretation says this. Occupying an institutional role does not make your personal reasons "disappear." But they are reasons for some actions and not others—for personal actions, not institutional ones. Someone can occupy several reason-giving relationships at once: family and institutional relationships, for example. These can conflict—as in example (*a*). There, facing question (1), you must decide between exercising two different kinds of agency—between performing a personal action or an institutional one. You have a personal reason (your family's need) for pursuing the business opportunity and an institutional reason (given by your responsibilities as scholarship officer) for appointing the best candidate instead. Here, the reason for performing the personal action could decisively outweigh the reason for performing the institutional one. However, in (*b*), question (2) concerns which institutional action to perform: it concerns how, not whether, to exercise institutional agency. Your personal reasons are irrelevant to that. You should do what the institution should do, through you.

The Persistence View, like mine, sees (1) as determined by the overall balance of personal and institutional reasons. But, unlike mine, it says the same about (2). It says that your family's need persists as an equally strong reason for appointing Smith. (It also says you should exclude that from deliberation, which you do.) So the problem for the Persistence

View is this: if it answers Yes to (1), how can it answer No to (2)? Since it says that the personal reason is equally strong in both cases, it will have to say that the contrary institutional reason is stronger in (b) than in (a). But it is hard to see how that could be so. In both cases, you occupy the same institutional role, with the same (implicit or explicit) instructions: so any institutional reason you have for appointing the strongest candidate in (b) is equally present in (a). It is true that in (b) but not (a) you (inadvertently) profit from the exercise of your institutional role; but the mere fact that a benefit happens to be causally downstream from an institutional action surely does not count against that action. So the Persistence View seems unable to explain why the two questions have different answers. It cannot account for the significance of whether a personal consideration falls within or outside the scope of institutional agency.

This challenge might seem easy to answer. There clearly is a difference between the two cases: the difference between helping your family instead of performing an institutional action and helping them as a result of your institutional action. So why can't the Persistence View just say that *that* is a difference between institutional reasons of different strengths? You have a stronger institutional reason against helping your family by performing an institutional action than you have against helping them instead of performing an institutional action. However, that just amounts to labeling the difference between the cases; it does not show how the Persistence View can explain its normative significance. What explains why that difference should be interpreted as a difference in the strengths of the reasons you have to comply with an institutional instruction? The role you occupy, and any associated institutional instructions, are the same in both cases.

To meet this challenge, the Persistence View needs somehow to find a reason for appointing the best candidate that is stronger once the institutional action has started in (b) than it is for starting it in (a). But attempts of that kind lead to the wrong distinctions. They either fail to distinguish (b) from (a), or they lead to implausible distinctions elsewhere. For example, it could plausibly be claimed that the reasons you have, in exercising your institutional role, to discharge your duties conscientiously, avoid corruption, act honestly, not subvert the institution, and so forth are stronger than the reasons you have to exercise the role at all. But there is no failure of any of those kinds in (b). Alternatively, one could try saying that you have a stronger reason to complete an institutional action well than to start it. That would distinguish (b) from (a). However, it would also distinguish your situation in (b) before and after you have started looking at the applications. And that is implausible: it is hard to believe that the proposition that you should appoint Smith is true before you have begun looking at the applications but becomes false once you have started.

My interpretation therefore seems to do a better job of explaining the difference between (a) and (b). It does so by explaining how context-undermining arises: your family's financial predicament is not relevant to the exercise of institutional agency. The Persistence View cannot give that explanation. The problem is to see how it can give another.

2. Return now to the commitment-based case. Here, the variations to consider are these. Suppose (a) my distress in testifying is unexpectedly severe—so severe that an expert psychologist assures you that I will suffer permanent harm unless you adjourn. Now it sounds as though you should break your promise. However, suppose (b) that we have discussed that extreme scenario in advance, at length, and I clear-mindedly and explicitly secured from you a promise not to take even distress of this severity into account in your decision making. Then, plausibly, you should keep it. Does that mean the promise has become inviolable—that it should be kept, no matter what? No: there is always some possible extremity that could justify breaking a promise, including this one. Suppose (c) that after you make the same emphatic promise, a deranged person credibly threatens to inflict permanent psychological harm on some other innocent person unless you break your promise by adjourning the hearing as soon as I show the slightest distress. That could give you a decisive reason to break it.

How do we account for the differences between these cases? We can start with the difference between (a) and (b). Here, two rival explanations are available. The Persistence View can say that in (a), the persisting reason you have to protect my welfare is strong enough to outweigh the promise, whereas the more serious promise in (b) carries more weight, making the reason to keep it decisive. My interpretation gives a different explanation. It sees the important issue as being whether my distress falls within the intended scope of a valid promise. To assess that, we need to ask two things: Is distress this severe reasonably regarded as falling within the intended scope of the promise I asked you to make? And if so, was I in a fit state of mind to exact from you a promise with that scope? In (a), the magnitude of the harm makes this dubious. In (b), however, the additional details are significant because they assure us that the answer to those questions is Yes. So (b) but not (a) is a case of context-undermining. Your severe distress is a reason for adjourning in (a) but not in (b).<sup>40</sup>

Thus, both views seem able to explain the difference between (a) and (b). But what about the difference between (b) and (c)? On my in-

40. My interpretation implies that there is no stronger reason for the action of adjourning in (b) than if my distress had only been slight. But it can allow that when the distress is greater, there is a stronger reason for (i) doubting that it is covered by your promise, (ii) wishing you hadn't made it, and (iii) trying to get you to rescind it.

terpretation, the difference again concerns whether the harm in question falls within the intended scope of a valid promise. This is true of the harm to me in (b), as we just saw. But it is not true of the harm to the other person in (c). Your promise, after all, is a promise to me: the reason you have to keep it is a reason to recognize my authority over whether *my* welfare is available as a reason bearing on your action. I have no such authority over anyone else's welfare—and no promise made by you could confer it.<sup>41</sup> So in (c), the harm to the other person is a reason for breaking your promise, and if the harm is great enough that reason could be decisive. In this respect, your promise is no different from any other: if you have promised to look after my scarf but then need to use it as a tourniquet to save someone's life, you should break your promise. In (b), by contrast, the harm to me is not a reason for breaking the promise since, in falling within the scope of the promise, it is subject to context-undermining.

The Persistence View lacks that explanation of the difference between (b) and (c). Both are cases in which there is a choice between keeping a serious promise and sparing someone permanent psychological harm. The Persistence View sees the harm as an equally strong reason for breaking the promise in either case. So it lacks the resources to explain the significance of whether the harm will befall the promisee or someone else. It is blind to the fact that I can have authority over my own welfare, but not that of other people.

It might seem tempting to reply that, since I have no authority to require you to ignore other people's welfare, you break no valid promise in (c). But that view of promises cannot be right. Whenever there is a decisive justification for breaking a promise, the promisee will lack authority to require the promisor to ignore that justification. But it is only because the promise remains in force that a justification is needed for breaking it. When I get you to promise to look after my scarf, I am not extracting from you a commitment to refuse to use it to save a life; but that does not prevent the life-saving action from being the (justified) breaking of a promise. On the contrary: it is because the promise is one there is a significant reason to keep that you owe me a strong reason for breaking it.

So the difference between (b) and (c) is not that the reason to keep the promise in (c) is weaker; it is that the strength of the reason that someone's distress gives for breaking it depends on whether it falls within the promise's scope. The problem here is structurally parallel to

41. You could say the words "I promise to grant you authority over whether Fred's welfare is a reason for me," but that would not succeed in creating any such authority—any more than saying "I bestow ownership of the *Mona Lisa* on you" would make you the owner of the *Mona Lisa*.

the one we found in the role-based case. There, the Persistence View cannot account for the significance of whether a personal consideration falls within or outside the scope of institutional agency; here, it cannot account for the significance of whether a harm falls within or outside the scope of your promise.<sup>42</sup>

There is a second problem with the Persistence View's treatment of the commitment-based cases. Tying together verdicts about action and deliberation, it says that whether you should exclude my distress from your deliberation depends on whether you should keep your promise. I say instead that it depends on context-undermining, which is determined by whether your distress falls within the scope of a valid promise, not whether the promise should be kept. To see the implications of this difference, look again at (c), where the promise should be broken to avert the threat to the other person. Suppose, accordingly, that you break it: at the first sign of distress from me you adjourn the hearing. If I ask you why, you can give me a good reason: it was necessary in order to avert a threat to someone else. But you cannot add, as a further part of your justification, that it was also necessary to spare me distress. My interpretation explains this: the promise is still valid, and my distress falls within its scope, so it remains unavailable as a reason for breaking it.<sup>43</sup> Although another person's welfare (over which I have no authority) can give you a reason for breaking your promise, my authority over the availability of my own welfare as a reason for you is unaffected. The Persistence View lacks an explanation of the continuing exclusion. It insists that my welfare *is* a reason for breaking the promise, and the case for excluding it from deliberation depends on whether you should keep the promise or not. On the Persistence View, there are two separate reasons for breaking the promise, and since you should break it, there is no explanation of why either of them should be excluded from deliberation.

42. Suppose what the promisee actually says is, "Promise me you will not let your reason to protect my welfare have an influence on your conduct of the hearing." Now the claim that you do have a reason to protect my welfare is built into the content of the promise. So how could the promise succeed in making that claim false? Doesn't it then become self-defeating? This kind of example, in my view, is one in which a promise can remain valid despite the promisor's erroneous belief. If I say, "I promise to give you my Louis XIV chair," unaware that it is really a Louis XV chair, then provided the intended reference is clear the promise remains in force. The promise about the hearing invites the same treatment. If, as I claim, the Persistence View is wrong, then a Persistence-View-presupposing speech-act of promising involves a false belief; but that need not invalidate the promise. A parallel treatment is available for institutional rules that are formulated in a Persistence-View-presupposing way.

43. You promised not to take account of my distress. So you are breaking your promise if your conduct of the hearing is determined by your sensitivity to my distress. But the reason you have for being sensitive to my distress is to avert the threat to the other person, not to spare me distress.

Can this objection be avoided by discriminating more finely between the actions on which these reasons bear? Breaking-your-promise-to-protect-the-threatened-person is a different action from breaking-your-promise-to-spare-me-distress. In (c), it is the first of those actions that is justified, and my distress is not a reason for *that*. It may seem that that is all the Persistence View needs to say in order to explain the deliberative exclusion.

But that is unsuccessful. According to the Persistence View, my distress persists as a reason for you to break your promise. So in (c) it has to say that there are two separate reasons to break it: my distress and the threat to the other person. They both therefore count in favor of this finely discriminated action: breaking-your-promise-to-spare-me-distress-and-protect-the-threatened-person. Since you should break the promise, and the Persistence View holds that the deliberative exclusion depends on whether you should keep it, it cannot explain why you should not break it in that way.

3. Last, we can turn to some variations on the meaning-based case. Suppose we increased the magnitude of Bill's enjoyment. Surely that would not change the verdict in this case. Even if (a) this train trip will be the most enjoyable experience in Bill's life, it would still not provide a basis for deciding which child's life to save. However, if we introduced considerations of other kinds, variations in weight could be significant. For example, could the fact that it will further someone else's interests justify you in deciding to take Bill rather than Ben? It depends on the magnitude of the interests at stake. If (b) this would merely spare someone else some inconvenience, No; but if (c) it would save someone else's life as well, then Yes—you could sensibly decide to take Bill rather than Ben for that reason.

On my interpretation, the explanation of this pattern of judgments is that the meaning of my life-saving decision restricts the reasons that bear on it to those relevant to answering the question, "How can you properly respect the value of your sons' lives?" Choosing between them by citing (a) Bill's enjoyment or (b) someone else's convenience fails that test. But citing (c) the loss of someone else's life meets it. A decision made on this basis can be properly respectful of the value of your sons' lives, in recognizing that the preciousness of their lives is a preciousness they share with others.<sup>44</sup>

Can the Persistence View reach the same set of verdicts about these cases? Apparently, it can do so as follows. The reasons to save your children's lives are incommensurate, not equal in strength, so the fact that Bill will enjoy it can give you more reason to take him on the life-

44. I think that is consistent with holding that you have a stronger reason to save the life of your child than that of a stranger.

saving trip than you would otherwise have without being a sufficient reason for taking Bill rather than Ben.<sup>45</sup> This is true no matter how much enjoyment is at stake; but the saving of someone else's life, by contrast, could give you a sufficient (perhaps, a decisive) reason for choosing between them. In this way, the Persistence View can produce sensible verdicts about action in these cases. It also needs to produce sensible verdicts about deliberation. You cannot properly invoke Bill's enjoyment as a reason for taking him on the life-saving trip, even when there are other good reasons for taking him as in (c). To explain this, the Persistence View can simply say that invoking reasons of enjoyment in deciding between your sons would still trivialize the decision—as surely it would.

But let us look at this more closely. We are comparing two rival answers to the question: "Why shouldn't Bill's enjoyment feature in deliberation about which son to take?" My proposed answer is: because, given the meaning of the decision, this is not a reason for taking him. The Persistence View denies that, saying instead that, even though this is a reason for taking him, it would trivialize my decision to invoke it in deliberation. We need to press this further. If it really were a reason, how would invoking it trivialize the decision?

An obvious-looking answer is this: although this is a reason, it is too small to be sensibly invoked in deliberation. After all, that seems the right thing to say about the opening example of the dinner guest, which the Persistence View assimilates to this one. However, there is a crucial difference. In the case of the dinner guest, there is a dominant reason that makes it overwhelmingly obvious what you should do. You should not attend to petty reasons when far more important ones make it obvious what you should do. But the train trip example has a different structure. You face a choice in which there are momentous reasons telling in opposite directions, and it is not obvious what to do. Given this, you are asking whether there are any further reasons for deciding one way or the other. And in a situation of this kind, the fact that a further reason is relatively small (compared to the momentous reasons that create the conflict) is not sufficient to show that it should be excluded from deliberation. For example: if there is only a small difference in age between your sons, or in their likelihood of survival, it could make sense to take that seriously as a basis for deciding between them.<sup>46</sup> My inter-

45. See Henry S. Richardson, *Practical Reasoning about Final Ends* (Cambridge: Cambridge University Press, 1994), 115–17.

46. I do not assert that these would be decisive reasons for choosing one child over the other. As I admitted in Sec. III, I don't know what verdict is right in this case. My claim is that it makes sense to take these seriously as reasons for choosing one child. That is consistent with thinking that there are stronger reasons not to choose either in preference to the other.

pretation explains this: that can be taken seriously, because it is consistent with respecting the value of your son's lives.

"Trivial," then, cannot simply mean small. Instead, one might suggest, it means frivolous, irrelevant. Small differences in the likelihood of survival could be seriously invoked as having a bearing on your decision: enjoyment on a birthday, whatever its magnitude, could not. But while that sounds sensible, to say that Bill's enjoyment is irrelevant to my decision concedes a context-undermining interpretation of the case. The irrelevance of enjoyment-facts is their failure to bear as reasons on a decision of this kind. Considerations of age or likelihood of survival do have that bearing, to some degree, because they bear on the question of how you can respect the value of your sons' lives; considerations of enjoyment do not, at all.

So the Persistence View cannot say that either. Again, it is left with a problem. It lacks a satisfactory explanation of how, if Bill's enjoyment really were a reason for taking him, invoking it would trivialize my decision. So it lacks an explanation of why this should be excluded from my deliberation.

Thus, in all three cases, the Persistence View has turned out to be problematic. It seemed at first to account neatly for our three examples; but having considered a further range of variations on those examples, we find that it struggles to account for the full pattern of verdicts that apply across that range—verdicts about both action and deliberation. It is not just that there are differences in the verdictive judgments implied by the two views. In each case, the context-undermining interpretation provides an explanation of the normative structure of the example that is a basis for confidence in those judgments. My interpretation explains the plausibility of the pattern of verdicts about both action and deliberation across the range of variations in the examples we have examined; the Persistence View is unable to do that.

## VI. ADDRESSING THE CHALLENGES

We began with a set of strong-looking challenges to accepting the existence of context-undermining. There was the core problem of seeing how an argument for context-undermining was even possible, given that showing why a consideration should be excluded from deliberation is not good enough; and on top of this there were three counterarguments—from explanatory simplicity, from doing what you should not, and from variation in weights. These can now be met.

How is it possible to argue for context-undermining? A two-stage response has been offered. First (in Sec. III), I explained what makes it true that there is context-undermining in the three target cases. According to that explanation, the status of a fact as a practical reason can

depend on other, contextual facts because of their effect on the nature of the action concerned. Specifically, this status can be affected by the kind of agency being exercised, the meaning of the decision being made, and the normative powers we recognize in respecting others. That gave us three answers to the question:

- (1) Why does *A* undermine *B*'s status as a reason?

The second stage, just completed in Section V, has involved vindicating those answers against the rival Persistence View. There, the aim was not to expand my answers to (1). It was to answer the different question:

- (2) Why should we accept my answers to (1)?

Answers to (1) are determinative reasons: reasons why certain facts lose the status of reasons. Answers to (2) are response-reasons: reasons for the response of accepting my answers to (1), in preference to the rival Persistence View.

To answer (2), I argued that my own interpretation makes better sense of the way in which verdicts about action and deliberation vary across a range of variations in our three examples. An argument of this kind presupposes that warranted claims about our reasons must cohere with our view of the ways in which it is rational to respond to those reasons in thought and action. Our access to truths about reasons comes from their impact on the practical thought and action we should endorse as rational (how else?); and it is by considering what best makes sense of that thought and action in its entirety that we will arrive at the most credible view of what those reasons really are and how they interact. This shows us how it is possible to argue for context-undermining, not just deliberative exclusion. The acceptance of context-undermining can be part of a view that best explains the full pattern of proper responsiveness to our reasons.

If this is right, not only can the core problem be addressed, but the three counterarguments fail too. According to the argument from explanatory simplicity, given any normative verdict for which a context-undermining explanation is offered, there is a simpler, outweighing explanation of the same verdict. We have found that that is wrong. The two rival views do not yield the same verdicts. The Persistence View is theoretically simpler, but its dual-verdict structure makes it too simple. It is unable to explain the normative significance of features that a context-undermining interpretation of the cases can account for.

The second counterargument was that in these cases there is a persisting reason for doing what you should not. This fails in a different way: it relies on one of two fallacious inferences. One is an inference from rationality to reasons: "Nepotistic action is not completely irrational; therefore there is a reason for it." This overlooks the possibility of

being rationally mistaken about your reasons. It can be rational to think that you have a reason and to act accordingly (to think that the building is on fire and therefore to try to escape), even if there is none.<sup>47</sup> So the inference is in general fallacious; and the argument for context-undermining gives us a further range of cases in which it fails. Personal loyalties are reasons for action, but not for institutional action. If so, those who believe that there is a reason for the nepotistic actions to which this explanation applies are mistaken. But that does not imply that their nepotistic actions are irrational.<sup>48</sup>

The other fallacious inference is: “Your nepotistic action will produce a good consequence: the help you give your relative. Therefore there is a reason to perform it.” This fails because the fact that an action produces a good consequence need not be a reason to perform it.<sup>49</sup> Perhaps if I hit you, you will behave with great dignity and poise. That is a good feature of the state of affairs that will result from hitting you, but not a reason to do it. The defense of context-undermining provides a further ground for rejecting this inference.<sup>50</sup>

The third counterargument concerned variation in weights: if amplifying a consideration can make it into a decisive reason, that suggests it is also a reason when it is smaller—just a less weighty one. But we have already seen that this argument fails in two of our three cases. Amplifying Bill’s enjoyment does not make it a decisive reason. And while amplifying my distress can sometimes justify breaking your promise, we saw that that is best explained by reference to whether the distress falls within the scope of a valid promise. Moreover, a similar reply is available in the role-based case. Yes, if a deranged person threatens to kill your

47. Compare Derek Parfit, “Rationality and Reasons,” in *Exploring Practical Philosophy: From Action to Values*, ed. Dan Egonsson, Jonas Josefsson, Björn Petersson, and Toni Rønnow-Rasmussen (Aldershot: Ashgate, 2001), 17–39, at 25, 29; Scanlon, *What We Owe to Each Other*, 22–32. They hold the less qualified view that, as long as you believe that there is a good reason for you to perform a given action, performing that action is rational (even if the belief is not).

48. I do think that, in a case in which I—an advocate of context-undermining—believe there is no reason to act nepotistically it would be irrational for me to do so.

49. To deny this is not to reject other, more plausible principles—such as the principle that, necessarily, something is good only if there is a reason to respond favorably toward it, for some favorable response; or the consequentialist principle that there is always most reason to do what produces the most good.

50. A third tempting inference: “Surely helping my relative get a job is something I have a reason to do: if I hear about an advertised job in her field, I have a reason to tell her. Nepotism is therefore a means to doing something I have a reason to do; so it must also be something I have *a* reason to engage in.” But here also, the principle that any means to an end you have a reason to achieve is a means you have a reason to take should be rejected. There may be ends you have a reason to achieve by some means and not others. If you have a reason to remove the pain in your little finger, that does not mean you have a reason to destroy the whole world.

child unless you appoint the weakest candidate on the job shortlist, you should do that. But my interpretation allows for this. It says that what you should do in acting as an agent of the institution is what it should do.<sup>51</sup> And universities should not (and surely, do not) require their employees to resist a threat like this, at the cost of their children's lives. Your situation in this case is like that of a bank teller in a hold-up. Banks should not and do not require their tellers to refuse to hand money to bank robbers.

Indeed, the emphasis in my discussion has been on showing that this kind of argument actually cuts the other way. It is the context-undermining approach that gives a better explanation of the significance of variations in the factors involved in these examples, and the Persistence View that is too crude.

## VII. CONCLUSION—A TEST?

How conclusive is this argument? Not completely, any more than normative argument ever is. For one thing, I have only defended my context-undermining interpretation of the examples against one rival: the Persistence View. But while that is a limitation on the argument, it is not a defect. As far as I can see, no one argument is powerful enough to defend a normative view against every rival at once. Here, as elsewhere, normative argument is a painstaking business of pairwise comparison. The challenge I set out to address was that claimed instances of context-undermining will always be more plausibly read as second-order response-reasons. It is not an objection to my argument that although it meets that challenge it goes no further.

The other respect in which obvious challenges have gone unanswered comes from the assumptions I have made. I acknowledged six of these early on but then made others subsequently. In particular, my argument against the Persistence View relied on assumptions about the correct verdicts in the range of cases I discussed—verdicts about both action and deliberation. I am not taking those judgments to be self-evident, and they could certainly be contested. Since our judgments about reasons and about what counts as rational reflection need ultimately to be justified together as part of the most coherent overall view, the possibility remains that the claims from which I have argued must in the end be abandoned. That possibility is not ruled out (it never is, in normative argument). Moreover, even granting the correctness of those verdicts, my discussion does not prove that the Persistence View is false.

51. Moreover, my view allows for the existence of cases in which you should use your role within an institution for ends other than its own. Those would be cases of justified theft: theft of another's agency.

It presents a challenge: to give a plausible explanation of those verdicts. I have set out the obstacles to doing that but cannot claim to have proved that it is impossible. The onus now shifts onto the opponent of context-undermining to answer this challenge.

Context-undermining, I have argued, comes from at least three different sources. I think there are others as well; but that could only be established by setting out further arguments of the form presented above. We opened with the question: what is the relevant difference between cases in which a persisting reason should be excluded from deliberation and cases of bona fide context-undermining? It would be nice to conclude with a neat test to distinguish the two. However, as far as I can see, the only way to do that is by attempting arguments of the kind I have offered here and seeing whether they work. For any proposed case of context-undermining, we must ask whether there is a cogent description of its structure, a demonstration of the ways in which alternative interpretations of the case are mistaken, and a satisfactory reply to the challenges described above.

We at least have a guideline to work from in looking for other cases. In the three discussed here, the context-underminer works by changing the nature of an action so that another fact loses its normative relevance to that action. A good way to look for other cases of context-undermining would be to think about the other ways in which that could be true.

Why does this issue matter? For two reasons. One is that we need to understand the relationships between reasons if we are really to understand why we ought to do some things and not to do others. Even if we agreed in all of our normative verdicts, practical philosophy would still have that elucidatory task. The other is that, as we have seen, the existence of context-undermining does actually make a practical difference: a difference in our verdicts about what is right, and a difference in how we should think about what to do.