

Hypothetical Consent and the Value(s) of Autonomy*

David Enoch

Hypothetical consent is puzzling. On the one hand, it seems to make a moral difference across a wide range of cases. On the other hand, there seem to be principled reasons to think that it cannot. In this article I put forward reasonably precise formulations of these general suspicions regarding hypothetical consent; I draw several distinctions regarding the ways in which hypothetical consent may make a moral difference; I distinguish between two autonomy-related concerns, nonalienation and sovereignty; and, utilizing these distinctions, I show that—and in a preliminary way, when—the objections to the moral significance of hypothetical consent fail.

I. INTRODUCTION (AND SOME EXAMPLES)

A patient arrives at your emergency room, unconscious. A blood transfusion will save her life. A blood transfusion is the kind of treatment that usually requires consent—without consent, it is usually morally impermissible to administer a blood transfusion. And the patient in front of you is not giving her consent. Of course, she cannot—she’s unconscious. But

* I presented earlier versions of this article in philosophy departments and law schools at Rutgers, Toronto, Ohio, the Hebrew University, Cambridge, Michigan, Harvard, UNAM, and University of Pennsylvania. I thank the audiences for the helpful discussions. For comments and conversations I especially thank Dani Attas, Selim Berker, Sarah Buss, Meir Buzaglo, Daniel Cook, Chris Cowie, Leora Dahan Katz, Tom Dougherty, Ruti Gavison, Daniel Goldstick, Kerah Gordon-Solmon, Abner Greene, Alex Guerrero, Scott Hershovits, David Heyd, Tom Hurka, Eric MacGilvray, Damian Melamedoff, Tristram McPherson, Sarah Moss, Peter Railton, Ram Rivlin, Amelie Rorty, Giorgio Sbardolini, Andrew Sepielli, Sergio Tenenbaum, Piers Turner, Peter Vallentyne, Alec Walen, Ben Zipursky, and several readers and editors for *Ethics*. My research was supported by the Israel Science Foundation, Grant 439/15.

Ethics 128 (October 2017): 6–36

© 2017 by The University of Chicago. All rights reserved. 0014-1704/2017/12801-0002\$10.00

perhaps you can still administer the life-saving treatment, for surely, had she been conscious, she would have given her consent. And perhaps this is enough to render the treatment morally permissible.¹

It is really hard to reconcile liberty and authority. How are we, if we are free, subject to the supposedly legitimate authority of the state? If we are born free, what can possibly make it legitimate that we are everywhere in chains? The natural thought is that consent would do just that. Being subjected to the authority of the state does not seem contrary to our liberty if we are only subject to it because we've consented to being subject to it. Alas, many of us have never given our consent. But perhaps the authority of the state can still be justified, for surely, had we been rational, we would have given our consent. And perhaps this is enough to render the state's authority legitimate.²

You want to cross a bridge. Unbeknownst to you, it is unstable, and if you get on it, you face great danger. For some reason, I can't convey this information to you sufficiently quickly (perhaps we don't share a language, and it will take time for the interpreter to arrive). I physically restrain you. This is the kind of thing that typically needs consent to be morally permissible. But you don't give your consent. In fact, you actively withhold consent. Still, perhaps my physically restraining you can nonetheless be morally permissible, for surely, had you realized the state of the bridge, you would have given your consent, or would have changed your plans to cross it. And perhaps this is enough to render the restraining morally permissible.³

In these examples—and in many more⁴—talk of hypothetical consent comes very naturally to us. The examples differ in important ways—the relevant hypothetical conditions, for one thing, are importantly different—but they all invoke hypothetical consent. Indeed, at least at first it seems hard to see how we could say all we want to say without resorting to hypothetical consent. But—as the literature on hypothetical consent shows clearly—the thought that hypothetical consent can make any normative difference whatsoever is, on reflection, perplexing.

1. This is a common example. See, e.g., Jeremy Waldron, "Theoretical Foundations of Liberalism," *Philosophical Quarterly* 37 (1987): 127–50, 139; Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, MA: Harvard University Press, 1990), 187. Below I'll be discussing several variations of this case.

2. See, e.g., David M. Estlund, *Democratic Authority: A Philosophical Framework* (Princeton, NJ: Princeton University Press, 2008), chap. 7, and the references therein. Note, however, that Estlund's own version of a hypothetical consent (the one he calls "normative consent") is in important respects atypical.

3. This example is loosely based on one in John Stuart Mill, *On Liberty* (London: Longman, Roberts & Green, 1869), chap. 5 (available at www.bartleby.com/130/). Mill does not, as far as I know, invoke thoughts about hypothetical consent here.

4. See, e.g., Arthur Kuflik, "Hypothetical Consent," in *The Ethics of Consent: Theory and Practice*, ed. Franklin Miller and Alan Wertheimer (Oxford: Oxford University Press, 2009), 131–61, 131–33.

I present the two main problems for the normative significance of hypothetical consent that I find in the literature in the next section: these are, first, that hypothetical consent can never substitute for actual consent, and second, that hypothetical consent always drops out of the explanatory normative story, because what does the relevant normative work is that in virtue of which, in the appropriate hypothetical conditions, consent would have been given. But I also insist—on the strength of intuitive examples, mostly—that even under the pressure of those objections, we should not conclude too quickly that hypothetical consent never matters. In order to offer a qualified defense of hypothetical consent, I present (in Sec. III) some crucial distinctions and then (in Secs. IV and V) proceed to show how—and when—the objections to hypothetical consent can be met. By the end of Section V, then, we will have some initial framework-setting conditions—what has to be the case for hypothetical consent to be possibly normatively relevant. But we will not yet have anything by way of a substantive account—a substantive answer to the question when, and why, hypothetical consent matters. I try to make progress on these questions in the following sections: in Section VI, I draw on ideas of the depth of the relevant commitments, centrality to the self, and Frankfurtean endorsement; in Section VII, I respond to an objection, namely, the worry that emphasizing the role of the agent’s deep commitments renders hypothetical consent normatively superfluous again; and in Section VIII, I distinguish two autonomy concerns that actual consent sometimes answers to, nonalienation and sovereignty, arguing that hypothetical consent can answer to the former but not the latter. (This distinction, I hope, is of interest also more broadly, not just in the context of discussions of hypothetical consent.)

II. TWO WORRIES

If we find it so natural—both in everyday and in philosophical contexts—to refer to hypothetical consent (at least when actual consent cannot be secured), why the suspicion that hypothetical consent is not, after all, morally significant?

I think that two central worries are relevant here.

A. *No Substitute for the Real Thing*

The first objection is straightforward enough.⁵ Consent has normative force.⁶ It can render impermissible things permissible (like your touch-

5. Stark refers to it as “Standard Indictment”; see Cynthia A. Stark, “Hypothetical Consent and Justification,” *Journal of Philosophy* 97 (2000): 313–34.

6. In Heidi Hurd’s memorable phrase, “consent turns a rape into love-making, a kidnapping into a Sunday drive, a battery into a football tackle, a theft into a gift, and a tres-

ing me, or your using my property). Consent—or its close relatives—can create duties (as in contracts). Giving consent is exercising a normative power.

But in these respects, a hypothetical consent just won't do the trick at all. If you complain about my using your property, and I note that you've consented (under the suitable conditions, etc.), then that is the end of the matter—a part of the bundle of rights that your property rights include are precisely the power, by your consent, to make it the case that it's permissible for me to use it.⁷ If you haven't consented, though, you haven't exercised that power of yours. How does it matter, then, that you would have, under some hypothetical conditions? It is of the very nature of normative powers—indeed, of powers more generally—that in order to have their standard (normative) effect, they must be exercised. A hypothetical consent, we may conclude, is not a pale form of an actual consent; it is no consent at all.⁸

Understood in this way, the thought that hypothetical consent can substitute for consent doesn't sound more plausible than the thought that if you're thirsty and there's no water around, it may be good enough that there would have been water, under suitably described hypothetical conditions, or the thought that you should get the medal even though you didn't win, because you would have won, under suitably described hypothetical conditions.⁹

B. *No Real Work*

Take the unconscious patient again. True, she would have consented to the blood transfusion, had she been conscious. But we can go deeper. Why would she have consented? Presumably, because of things like that she wants to go on living; that she's afraid of death; that she cares intrinsically about some projects, relationships, people—things of value in her life, which she can continue engaging in hopefully constructive, valuable ways if she survives, but not otherwise. But all these things—aren't they

pass into a dinner party." See Heidi Hurd, "The Moral Magic of Consent," *Legal Theory* 2 (1996): 121–46.

7. Of course, I'm simplifying here about property rights, the precise content of which—I'm sure—heavily depends on context. Such complications won't matter for our purposes.

8. This paraphrases Dworkin's famous criticism of Rawls's use of hypothetical contract. See Ronald Dworkin, "The Original Position," in *Reading Rawls*, ed. Norman Daniels (New York: Basic, 1975), 17–52, 18. For the record, I'm not sure it's a fair criticism of Rawls, but the details are irrelevant for the discussion here.

9. Actually, this last example is not entirely clean. Think about someone who comes in second, and then it's found out that the competitor who came in first was using forbidden performance-enhancing drugs. You may want to say that the one who came in second actually won. But at least as natural a description would be that he hasn't, but would have, if it weren't for the cheating.

reasons enough for you to administer the blood transfusion? The hypothetical consent seems not to do any normative work anymore; rather, the underlying normative reasons that make it the case that the patient would have consented (had she been conscious) also make it the case—on their own—that you should administer the blood transfusion. The hypothetical consent then becomes not what makes it the case that it's permissible to administer the blood transfusion, but a by-product of the factors that make it the case that it's permissible to do so.

Or take the political example again. If someone has not given his consent to the authority of the state—indeed, is actively refusing to give his consent—but would have consented, had he acted rationally, there must be something in virtue of which he would have consented (that there must be such reasons seems to follow from the thought that he would have been rational to give his consent). Perhaps, had he been acting rationally, he would have given his consent because the alternative to an authoritative state is a terrible Hobbesian state of nature, or perhaps it's because a democratic state's authority is legitimate as the only public order in which people treat each other as free and equal, or some such. But then, if some such considerations make it the case that the citizen would have consented had he acted rationally, they seem to justify the authority or legitimacy of the state directly. Once again, then, the hypothetical consent drops out of the normative picture. Rather than legitimize the state (or some such), it is merely the normative by-product of that which legitimizes the state.

Perhaps we can make this line of thought more precise, in terms of what I will call the Transitivity Argument.¹⁰ Let “Normative Upshot” designate whatever it is that hypothetical consent may be relevant to—the permissibility of administering the blood transfusion, the legitimacy of the state, or some such. And let “Underlying Reasons” stand for whatever considerations make it the case that, in the relevant hypothetical conditions, consent would have been given. Then we get the following:

1. The Normative Upshot holds in virtue of hypothetical consent.
2. Hypothetical consent would be given in virtue of the Underlying Reasons.¹¹

10. This line of thought is fairly common in the literature. See, e.g., Thomson, *Realm of Rights*, 187; Thomas Hill, “Hypothetical Consent in Kantian Constructivism,” *Social Philosophy and Policy* 18 (2001): 300–329, 326–27; Gopal Sreenivasan, “‘Oh, But You Should Have’: Estlund on Normative Consent,” *Iyyun* 58 (2009): 62–72, 68. The more precise argument that follows is, as far as I know, original, though I found something close in Don Hubin, “Taking Hypothetical Disagreement Seriously” (unpublished manuscript).

11. As an anonymous referee correctly pointed out, there's an ambiguity here: it's one thing to ask in virtue of what is it the case that “under hypothetical conditions C consent would have been given” and another to ask in virtue of what would agent A consent in hy-

3. The in-virtue-of relation is transitive.
4. Therefore, the Normative Upshot holds in virtue of the Underlying Reasons (from (1), (2), and (3)).
5. Therefore, hypothetical consent does no normative work in justifying the Normative Upshot (from (4)).¹²

I return to this argument in Section IV. For now, though, I want to quickly dismiss some initial objections and flag some further relevant issues.

Because Underlying Reasons have been stipulatively defined as that in virtue of which hypothetical consent is present, the only way to reject premise (2) seems to be to assert (in a specific case, perhaps) that there is nothing in virtue of which the relevant hypothetical consent is present, that it is groundless. Now, this is a possible move, of course—in-virtue-of chains, like many others, come to an end somewhere. But it seems like an especially implausible place to bring the chain to a stop.¹³ After all, we've chosen, among the infinitely many possible hypothetical situations, the one situation in which consent counts (for the relevant Normative Upshot). Why that one, though? Presumably, because in that hypothetical scenario—when you are conscious, when you are rational, when you are fully informed—you are more sensitive to the relevant considerations, that is, it seems, to the ones in virtue of which you would consent. So the hypothetical consent is not groundless.¹⁴

The transitivity of the in-virtue-of relation (premise (3)) may, I guess, be rejected. But at least without some story making such a move plausible, it would seem to be a high price to pay. Related suspicions, though, may be more productive here. In the grounding literature, some people distinguish kinds of grounding—say, metaphysical, causal, and normative grounding. If so, the question whether grounding is transitive across kinds of grounding becomes interesting: if A metaphysically grounds B, and B normatively grounds C, does it follow that A grounds C, and if so, what kind

pothetical conditions C. But in many cases the same answer will fit both questions, because often, the agent would, in hypothetical conditions C, consent in virtue of some fact, and because of that, “in conditions C A would consent” will be true in virtue of that same fact. For now, I am happy to restrict everything I say to just those cases in which this is true—I think that very often, the specification of the hypothetical conditions will guarantee that this is the case.

12. Notice that the conclusion is about hypothetical consent not doing any work in justifying the Normative Upshot. This is consistent with its doing work in explaining how it is that the Underlying Reasons justify the Normative Upshot. In other words, it is consistent with the role hypothetical consent is playing in the Transitivity Argument itself.

13. Hubin makes a related point; see Hubin, “Taking Hypothetical Disagreement Seriously.”

14. Hill, “Hypothetical Consent in Kantian Constructivism,” 305, makes a closely similar point.

of grounding relation is this?¹⁵ Similar questions may be asked directly about the in-virtue-of relation invoked in the Transitivity Argument.¹⁶ Perhaps in some applications of the argument, the in-virtue-of relation in premise (2) is causal, whereas the one in premise (1) is normative. This may be so, for instance, if the patient's (motivating) reasons why she would have consented had she been conscious are not in fact good (normative) reasons, namely, they don't count in favor of accepting the transfusion. In such a case, it's not clear that the transitivity move succeeds—it may be guilty of equivocation. Still, for the most part we can bypass this difficulty, I think, by focusing on those hypothetical conditions in which it seems like the hypothetical consent holds normatively in virtue of the Underlying Reasons. This is why the political example is important—the in-virtue-of relation that holds between the fact that someone would have consented had he been rational and his Underlying Reasons for consenting to the state's authority seems to be normative through and through. Still, the equivocation worry will return below.

Even if the argument is sound, its conclusion is consistent with hypothetical consent playing an epistemic role. Think, for instance, about using "What would Jesus do?" as a tool in one's practical thinking. Notice that this amounts to the use of a hypothetical test of sorts. But in order for such use to make sense, one doesn't have to believe that the fact that Jesus would do so-and-so makes it the case that so-and-so is the thing to do. One may think of the fact that Jesus would do so-and-so merely as evidence—perhaps conclusive evidence, but still evidence—that so-and-so is the thing to do. Analogously, then, even if the Transitivity Argument shows that hypothetical consent never grounds the relevant Normative Upshot, hypothetical consent may still serve as evidence for the Normative Upshot (if, that is, for some reason it's easier for the relevant agent to know whether there would be hypothetical consent than it is to consider the Underlying Reasons directly; perhaps this is atypical, but surely it's not impossible). I take it, though, that often (as in the examples I started with) we seem to want more for the significance of hypothetical consent, more than merely evidence for reasons that are independent of consent. True, there may be cases in which hypothetical consent (and, indeed, ac-

15. See Selim Berker, "The Unity of Grounding," *Mind* (forthcoming), and the references therein. Berker argues, on the basis of the transitivity of such (purported) structures, that the grounding relation is in fact unitary.

16. In previous versions of this article, I characterized the relation the transitivity of which is used in the argument as the grounding relation. But as two referees (independently) noted, it was not clear that I used that term in a way that was continuous with the increasingly large metaphysical literature on grounding. For my purposes, then, it is best to keep with the nontechnical language of "in virtue of" and to rely on the grounding literature, if at all, by analogy.

tual consent) plays just such an epistemic role,¹⁷ but in some cases it does more. Clearly, actual consent often does more, so allowing only an epistemic role to hypothetical consent will be to give up on much of the intended role of hypothetical consent.¹⁸

These rather quick attempts at resisting the Transitivity Argument, then, do not work.

C. Combining the Two Worries: Will or Reason?

As if these worries about hypothetical consent are not bad enough, when combined they can be more worrying still. For they give rise to a powerful dilemma.

Jeremy Waldron famously distinguishes, in the political context, between consent theories that are will based and those that are reason based. Will-based theories use consent to show that the relevant person's will is engaged in the right kind of way. Reason-based theories use consent to show that the relevant person has the right kind of reasons. We can use this distinction to offer the following dilemma: hypothetical consent can be normatively relevant either for will- or for reason-related reasons.¹⁹ If it's supposed to be relevant for will-related reasons, though, the objection in Section II.A seems especially powerful, for that I would have consented doesn't show that my will is engaged in the right kind of way, only that it would have been. If it's supposed to be relevant for reasons-related reasons, the Transitivity Argument applies especially powerfully, as then the Underlying Reasons seem to be doing all the relevant normative work. The challenge to defend some normative significance for hypothetical consent then becomes that of avoiding both pitfalls—we must show, roughly, that hypothetical consent is sufficiently about the actual will of the relevant person to avoid the Transitivity Argument, but also that it's about his or her will in a way that avoids the problem of consent not being a pale form of consent, but rather no consent at all.

III. ANOTHER EXAMPLE, AND TWO DISTINCTIONS

Despite the naturalness of resorting to hypothetical consent, then, there are reasons for serious suspicion that perhaps, at the end of the day, hy-

17. Tom Dougherty stressed to me that I was using high-stakes examples, and that perhaps with low-stakes examples our intuitions are different—he mentioned cases like it being permissible to gently touch a passenger on a train in order to wake them up at the last stop, based on the plausible thought that they would have consented. Perhaps in such cases, the significance of hypothetical consent is merely epistemic.

18. I thank Alex Guerrero and Sergio Tenenbaum for pressing me on related points. Furthermore, if hypothetical consent is going to have merely an epistemic role, as an indication of something else, then it must be asked, of what? The only plausible answer I can think of is in terms of deep commitments, which I discuss later in the text.

19. This is not what Waldron uses it for; see Waldron, "Theoretical Foundations of Liberalism."

pothetical consent cannot be normatively significant. It cannot do the work of actual consent, because for normative powers to make the difference they are supposed to make, they have to be exercised, and we have some strong initial reasons—captured by the Transitivity Argument—to suspect that the hypothetical consent itself doesn't do any normative work.

I am not willing to give up on hypothetical consent just yet, though. There are cases in which—despite the worries from the previous section—it's very hard not to assign normative significance to hypothetical consent. To see this, compare the original unconscious patient case to the following variation. In this variation, the unconscious patient arriving in your emergency room is a Christian Scientist and is therefore strongly committed, on religious grounds, to not allowing a blood transfusion to be administered to her. In terms of actual consent, there is no difference between the two patients—both of them are unconscious. We can also stipulate that it would be good for both to survive, perhaps even to the same degree.²⁰ The Underlying Reasons that both have to go on living are equally strong. Still, isn't it clear intuitively that there is a normatively relevant difference between the two cases? Perhaps the difference is not sufficiently potent to make it the case that it's permissible to administer the transfusion in the first case but not in the second. I'm not sure about this. But I am quite sure that there is at least some normative difference between the two cases. The natural way to capture the difference is by talk of hypothetical consent: the first patient would have consented, had she been conscious; the second one would not have, because of her Christian Scientist commitments. So it's worth our time to see whether some normative significance for hypothetical consent can be salvaged.

One important step in this direction, then, is to distinguish among different kinds of normative significance, different ways in which hypothetical consent may (be thought to) make a normative difference. In particular, if hypothetical consent makes a difference, we can ask what kind of difference, and to whom.

A. *Different Normative Upshots*

I already noted that consent and its relatives make all sorts of normative differences. My willfully entering the boxing ring with you arguably makes

20. Thomson (*Realm of Rights*, 189) suggests that in such a case, if we should not administer the transfusion to the Christian Scientist, this is because it would be bad for him (say, to survive in virtue of a measure he is strongly opposed to on religious grounds). This seems wrong to me: surely, we can come up with an example in which, all things considered, surviving even by those means would be good for the Christian Scientist. We can even come up with an example where the Christian Scientist himself will acknowledge as much. The nature of the religious prohibition is—or at least can be, for these are hypothetical examples we're dealing with—deontological rather than good related.

your punching me—usually, a morally impermissible action—into a permissible one. Here the Normative Upshot is that of permissibility. My promising to meet you for lunch creates for me a duty to be at the cafeteria at noon. Here, the Normative Upshot of the promise (importantly different from consent, but certainly in the same normative family) is that of obligation, or duty. Perhaps if I consent to the authority of an arbitrator, she now has the normative power to create duties for me. Perhaps here the Normative Upshot of my exercising the normative power of consent is that of creating further normative powers (in the arbitrator). Perhaps sometimes consent is not itself a reason, but rather it defeats a reason, or perhaps it defeats a defeater of a reason. Indeed, you may think that this is what's going on in patient cases—there are consent-independent reasons to save, nonconsent (or refusal, perhaps) defeats those reasons, and consent (perhaps including hypothetical consent) may defeat this defeater.²¹ And so on.

When asking about the normative significance of hypothetical consent, we should not be working with an overly poor menu of Normative Upshots.²² Even if, for instance, hypothetical consent—unlike actual consent—never makes an impermissible action permissible, this does not mean that it does not make a normative difference, because it may have other Normative Upshots.²³

In particular, two points are worth stressing here. First, normatively mattering can be a matter of degree. In the boxing example (with actual, though perhaps implicit, consent), for instance, you may think that the wrongness of the entire practice makes it the case that my consent does not in fact make your punching me morally permissible. Even if you do think that, though, I am pretty confident you don't think that my consent is normatively inert. You probably think that your punching me, while still wrong, is not as seriously wrong as it would have been had I not given my consent. Or perhaps the punching—while still wrong—is wrong in somewhat different ways, or for somewhat different reasons, than it would have been but for the consent. Similarly, then, it's possible that hypothetical consent does not make an impermissible action permissible but still makes

21. I owe this point to Kerah Gordon-Solmon.

22. I think, though, that we can rule out the possibility that hypothetical consent matters merely instrumentally. The problem is that mattering instrumentally is a causal matter, and hypothetical consent is not actual, and so presumably it is causally inert. (Of course, *taking* hypothetical consent to matter can matter instrumentally, and this may be an important consideration to take into account, e.g., in institutional settings. But this is a different matter.)

23. Kuflik ("Hypothetical Consent") correctly criticizes Thomson, noting that hypothetical consent may be normatively relevant even if there's no moral status for which it's either a necessary or a sufficient condition.

a normative difference (think again about the comparison between the two patient cases above).²⁴

Second, even if hypothetical consent can never do the normative job that actual consent does, this does not mean that it doesn't do any normative work at all. There are many ways for things to be normatively significant, and not being significant in the way that actual consent is just does not entail not being normatively significant at all.²⁵

B. *To Whom*

Forget talk of permissibility and impermissibility. Let's talk of reasons for action. Consent (and its relatives) can affect everyone's reasons for action. By my consent to your touching me, I can make it the case that reasons you had not to touch me are no longer in place. I can make it the case that the weight of other reasons of yours changes. By the exercises of related normative powers I can bring reasons for action into existence.²⁶ I can also affect my own reasons for actions. I can, for instance, create duties for myself, and so on.

It may be important to ask, assuming that hypothetical consent can make a difference in terms of reasons for actions, whose reasons it may affect. The relevant distinction here is, I think, that between my hypothetical consent affecting my own reasons for action and it affecting others' reasons for actions. You may think, for instance, that in order to affect others' reasons for action—especially in the way actual consent does, say, by rendering certain impermissible actions permissible—consent has to be actual, but that the fact of hypothetical consent may affect one's own reasons for action.²⁷ That is, you may think that the mere hypothetical consent of an agent never makes a difference to the reasons that apply to others' treatment of her, but that it can make a difference to her own rea-

24. Perhaps—I am really not sure about this—if something makes any normative difference, it just follows that we can describe a possible case in which it makes the difference between permissibility and impermissibility. Perhaps, for instance, if my consent to boxing makes a normative difference, then while it needn't make your punching me morally permissible, it makes impermissible some intervention by a third party that would have been permissible otherwise. If this is so, then making a normative difference anywhere entails making the difference between impermissibility and permissibility somewhere. This would require weakening some of the claims in the text here, but not, I think, in crucial ways: it would still remain true that in a given context hypothetical consent could make a normative difference without making—in that context—the difference between impermissibility and permissibility.

25. This is a point that Stark emphasizes in "Hypothetical Consent and Justification."

26. For requests, see David Enoch, "Giving Practical Reasons," *Philosopher's Imprint* 11 (2011), <https://quod.lib.umich.edu/cgi/p/pod/dod-idx/giving-practical-reasons.pdf?c=phimp;idno=3521354.0011.004>. For orders and authority, see David Enoch, "Authority and Reason-Giving," *Philosophy and Phenomenological Research* 89 (2014): 296–332.

27. This, I think, is Stark's view.

sons. Or you may want to insist that hypothetical consent does make this kind of a difference for others' reasons as well, or perhaps even only for others' reasons.²⁸

There are many ways, then, in which hypothetical consent may make a difference: It may make a difference by rendering impermissible actions permissible, but even if it does not, it may have other Normative Upshots. It may make a difference to the reasons for action of the person whose consent it is, or to others' reasons for action, or both. Still, for the most part in what follows I will be focusing on the narrower questions—whether hypothetical consent can make the kind of normative difference actual consent often makes, by affecting other people's reasons for action, and in particular, by rendering otherwise impermissible action permissible. This is the kind of case where hypothetical consent is most often invoked,²⁹ as well as the kind of normative significance that was relevant in the three examples with which I started. The distinctions between the different ways in which hypothetical consent can make a normative difference will nevertheless be of importance later on.

IV. THE TRANSITIVITY ARGUMENT, AND MULTIPLE REALIZABILITY

Recall the Transitivity Argument:

1. The Normative Upshot holds in virtue of hypothetical consent.
2. Hypothetical consent would be given in virtue of the Underlying Reasons.
3. The in-virtue-of relation is transitive.
4. Therefore, the Normative Upshot holds in virtue of the Underlying Reasons (from (1), (2), and (3)).
5. Therefore, hypothetical consent does no normative work in justifying the Normative Upshot (from (4)).

I want to grant the argument up to premise (4). What I want to question now is the move from premise (4) to premise (5). In other words, even granting the transitivity move (and so perhaps, following the discussion above in Sec. II.B, focusing on cases in which premise [2] is true with the

28. One way of understanding Stark's ("Hypothetical Consent and Justification") master argument is as starting off agreeing that one's hypothetical consent doesn't directly affect how one is permissibly treated by others; insisting that it nevertheless may have another kind of Normative Upshot, affecting one's own reasons for action; and then claiming that, in the political context, sometimes how one can be permissibly treated is affected by what reasons for action one has oneself.

29. Stark (*ibid.*) argues that in the political context, at least in Rawlsian views, hypothetical consent should be understood as relevant, but in a very different way, with a different Normative Upshot. For reasons that are irrelevant here, I am unconvinced.

“in virtue of” understood normatively), I want to show that the move to premise (5) can be resisted.

It will be helpful to start by thinking about analogous structures elsewhere—in particular, about the explanatory force of the special sciences.³⁰ Think, then, of the following argument:

1. Biological properties are instantiated in virtue of chemical ones being instantiated (for *reductio*).
2. Chemical properties are instantiated in virtue of physical ones being instantiated.
3. The in-virtue-of relation is transitive.
4. Therefore, biological properties are instantiated in virtue of physical properties being instantiated.
5. Therefore, chemical properties do not do any (causal? explanatory?) work re biological properties.

I take it that the conclusion is clearly false. Grounded, supervening properties can do genuine explanatory work. The most natural way of seeing that is focusing on multiply realizable, or multiply groundable, grounded properties. Perhaps the chemical property of solubility in water can be realized or grounded in different physical structures—P1, P2, and P3. If so, it’s quite possible that the best explanation of some biological fact is that some material is soluble in water, rather than the explanation in terms of its possessing the physical property P2. This is so if possessing P1 or P3 would have been equally good for explaining the relevant biological fact. In such a case, an explanation at the chemical level is more informative because it is at the right level of generality compared to its counterpart in purely physical terms. Thus, what seems to be doing the best explanatory work is the supervenient, grounded chemical property, rather than the underlying, subvening, grounding, physical property. Analogously, that moral properties supervene on and hold (normatively) in virtue of nonmoral ones does not undermine moral properties’ explanatory potency. Perhaps, for instance, what explains social instability is the relevant society’s being unjust. This can be so even if the injustice itself is grounded in some nonmoral property (say, about the distribution of resources, or of power), so long as there’s a range of nonmoral phenomena that could make it the case that the society was unjust, and as long as injustice thus grounded would also have led to instability.³¹ If

30. I heard David Estlund make a comment anticipating this move in response to an objection, in the context of defending his version of hypothetical consent, i.e., normative consent. I don’t think that he’s pursued this line in writing. I thank Arnon Levy for related discussion.

31. See Geoff Sayre-McCord, “Moral Theory and Explanatory Impotence,” *Midwest Studies in Philosophy* 12 (1988): 433–57, sec. 7; Peter Railton, “Moral Explanation and Moral

so, it's the moral explanation, not the nonmoral one, that is cast at the right level of generality for being the best explanation.

Can we say something analogous, then, in order to block the transition from premise (4) to premise (5) in the Transitivity Argument against the normative significance of hypothetical consent? In order to do so, we would have to show that hypothetical consent can hold in virtue of different Underlying Reasons, and that, thus grounded, hypothetical consent could still be that in virtue of which the Normative Upshot holds. What we need, in other words, are intuitively plausible cases where all you need to know in order to know that the Normative Upshot is there is that there is hypothetical consent; you needn't care about why it is that there is hypothetical consent. Notice that this is what we would say about the normative significance of actual consent: even if actual consent too is normatively grounded (say, in one's reasons for consent), often the Normative Upshot of consent is secured by the consent alone, regardless (within some constraints) of the reasons for consent.

So what we need is something like the following: Perhaps there are any number of potential reasons for which the unconscious patient could have consented, had she been conscious. Perhaps she could have done it out of fear of death, out of blind obedience to medical authority, out of *joie de vivre*, out of a sense of duty to one's loved ones, out of the desire to complete one's philosophy paper, and so on. If it's plausible to say that in (roughly) all and only cases in which the patient would have consented had she been conscious, it's permissible to administer the blood transfusion, regardless of why it is that she would have consented, then the Transitivity Argument fails—the move from premise (4) to premise (5) is blocked, and hypothetical consent, grounded though it is in other things, still does genuine normative work.

Notice that this response to the Transitivity Argument does not amount (yet) to a positive argument for the normative significance of hypothetical consent. It does not even amount to a rejection of the Transitivity Argument against such significance. What it does is open the door for such a defense. The defense itself will have to depend on the specific details of the relevant hypothetical consent theory, in particular, the specific hypothetical conditions invoked, and the specific Normative Upshot the theory is out to secure. Perhaps, for instance, hypothetical conditions that include things like "had she been conscious" (as in the patient case) or "had he been well informed" (as in the bridge case) are ones that satisfy this condition (because we don't care why he would have changed his plans re the bridge had he known about its state; we only care that he

Objectivity," *Philosophy and Phenomenological Research* 58 (1998): 175–82; Nick L. Sturgeon, "Thomson against Moral Explanations," *Philosophy and Phenomenological Research* 58 (1998): 199–206.

would do so). But perhaps hypothetical conditions that include “had he been rational” conditions (as in the political authority case) are in worse shape regarding the challenge posed by the Transitivity Argument, because perhaps in asking about what the citizen would have consented to had he been rational, we are already asking about the reasons for consent, rather than the consent itself.³²

The plausibility of this response to the Transitivity Argument can thus depend on the specific hypothetical conditions invoked. It can also depend on the relevant Normative Upshot, and here the distinctions from the previous section may be helpful. You may think that different Normative Upshots hold in virtue of the Underlying Reasons and in virtue of the hypothetical consent. You may think, for instance, that the reasons that make it the case that the unconscious patient would have consented to the blood transfusion had she been conscious are reasons for her,³³ but that the reason why it’s permissible for you to administer the transfusion is not her *joie de vivre* (or whatever), but the fact that, for whatever reason, she would have consented had she been conscious. Or perhaps her *joie de vivre* counts in favor of you administering the transfusion, but the only thing that makes it *permissible* for you to do so is that she would have consented.

Much remains, then, for the substantive, detail-specific discussion of specific suggestions as to the way in which hypothetical consent matters. This, it seems to me, is as it should be. Still, this section establishes two important results. First, it shows that the Transitivity Argument against the normative significance of hypothetical consent is by no means conclusive. And second, it helps to establish a desideratum for hypothetical consent theories—if you want to put forward such a theory, you had bet-

32. However, things here may depend on the specific conception of rationality employed. Under a purely instrumental conception of rationality, the point in the text does not seem to go through. Also, if the case can plausibly be made that what matters for the relevant purposes is not the reasons for consent, but rather that there are reasons for consent (whatever they are), then the general strategy of resisting the Transitivity Argument can again be utilized here as well, even under a substantive conception of rationality (as responsiveness to reasons).

Let me note quickly here, without the further discussion that would be needed to establish this point, that I think that Estlund’s “normative consent” version of a hypothetical consent theory is especially vulnerable to the point in the text here. Because his hypothetical conditions are normative, the Transitivity Argument is especially appealing. See here the exchange between Sreenivasan (“Oh, But You Should Have”) and David Estlund, “Reply to Commentators,” *Iyyun* 58 (2009): 73–88.

33. Or—following a point from Sec. II—perhaps her reasons for (hypothetically) consenting are not good reasons at all, perhaps they’re just her motivating reasons; perhaps, in other words, hypothetical consent is here present causally because, but not normatively because, of the Underlying Reasons.

ter be able to make the multiple-realizability kind of point,³⁴ with the details of your specific theory plugged in, in a substantively plausible way.

V. WHEN IS IDEALIZATION IN GENERAL ACCEPTABLE?

Recall the thought that hypothetical consent is not a pale form of consent; rather, it is no consent at all. There is a natural sense in which this is obviously right. But we should resist too quick and broad a generalization. Sometimes, after all, going hypothetical, or idealizing, is an acceptable move. And I think it will be helpful to have a more general look here and spend some time on the more general question—when doing theory, any kind of a theory, when is idealizing or going hypothetical acceptable?³⁵

A theory is offered, one that ties some phenomenon to our relevant responses. Perhaps, for instance, a theory is offered about the relations between (phenomenal) color and our color appearances, so that to be red is just to appear red. Or perhaps a theory of values is offered that ties them very closely to what we value, so that to be of value just is to be valued. But counterexamples immediately come up: Sometimes something appears red to me even though it isn't, and sometimes it doesn't even though it is. Sometimes people (even I) value things that aren't of value and fail to value things that are. A natural move then is to idealize: Perhaps to be red is not to appear red, but rather to appear red to normal observers, in good lighting conditions. Or perhaps to be of value isn't to be valued, but to be valued by the right people, in the right conditions. And of course, any such idealization can be thought of as hypotheticalization—perhaps being red is tied to how things would have looked to you, had you been a normal observer, in good lighting conditions. Perhaps being of value is tied to what you would have valued, had you been in the appropriate conditions.

Whether such idealization is respectable depends on what the underlying motivation was for going for the relevant view in the first place. Think again about the water example. Presumably, your reason for looking for water is that it would quench your thirst. Hypothetical water won't do that—it won't even go part of the way toward doing that. So in the absence of water, going for hypothetical water amounts to cheating—it is disconnected from the underlying motivations of looking for water and therefore offers no satisfaction, not even partially.

34. I haven't established—nor can I—that the only way of resisting the Transitivity Argument is by employing the line in the text. If there are others, the point in the text here should be qualified accordingly. Indeed, the point about different kinds of in-virtue-of relations (from Sec. II and from the previous footnote) may again be relevant.

35. In what follows I draw, in general, on David Enoch, "Why Idealize?," *Ethics* 115 (2005): 759–87, and also on some paragraphs in David Enoch, "Against Public Reason," *Oxford Studies in Political Philosophy* 1 (2015): 112–42.

What about colors? Is idealizing here a way of cheating, avoiding counterexamples in an ad hoc way? Or are ideal, hypothetical responses enough here? The answer depends on the philosophical motivations underlying the relevant account of colors. If they are all about actual observers and their actual responses, then going hypothetical (because of the pressure from counterexamples) is cheating. What counterexamples show us, in such a case, is not that we should settle for observations in hypothetical conditions, but that we should resist the attempt to tie colors and our appearances of colors as closely together as all that. But if the underlying motivations of such an account are consistent with settling for hypothetical conditions, then all is well, and the idealized response-dependence account may still be a good idea.

Similarly for values and what we value. If the underlying motivations for offering a response-dependence view of values are tied to actual people and their responses, then what we should do in the face of the obvious counterexamples is not idealize (in an ad hoc way), but rather reject response dependence altogether. But if the idealization can be motivated in a way that's consistent with going for a response-dependence view and with the philosophical motivations for so going, then all may still be well.

And so, we have a test for when idealization is a legitimate philosophical move. It is when it is motivated and, furthermore, when the offered motivation is consistent with the motivations for going for the initial, non-idealized view (the one that was devastated by obvious counterexamples). This is why hypothetical water is out, why (perhaps) some idealized response-dependence views of colors may be in, and why (as I argue elsewhere) idealized response-dependence views of normative concepts are out.³⁶

Let us end here our detour on idealization in general. Now we can note how this test applies to the case of hypothetical consent and, in particular, to the question when it can do the normative work that in more simple cases actual consent does. The thing to do is to ask why it is that actual consent matters, when in fact it does. And then we need to ask whether the concerns to which actual consent answers are also answerable by hypothetical consent (and if so, which hypothetical consent, in which hypothetical conditions).

So we need to think about actual consent. But already before doing that, we can note how surprising it will be if we get a general, context-insensitive answer to this question. It seems much more plausible to suppose that actual consent matters in different ways in different contexts,

36. Again, see Enoch, "Why Idealize?" But see also David Sobel, "Subjectivism and Idealization," *Ethics* 119 (2009): 336–52; David Enoch, "Idealizing Still Not Off the Hook: A Reply to David Sobel" (unpublished manuscript); Dale Dorsey, "Idealization and the Heart of Subjectivism," *Noûs* 51 (2017): 196–217.

for different purposes. So perhaps the way to progress is to acknowledge such complexity, to see how actual consent matters in different contexts, and then to apply the test from above to see whether (and what kind of) hypothetical consent can matter in similar ways. I do just that (in a somewhat preliminary way) in Section VIII. Before that, though, we need more examples, and we need to think about depth of commitments, centrality to the self, and endorsement.

VI. DEPTH

Consider now five patient cases. We've already described two of them. But we need all five now.

The first two patients are conscious, and they actively refuse to consent to being given a blood transfusion. But they do so for different reasons. The Conscious Christian Scientist does this because of her religious commitments. The Conscious Anxious Patient does this because of his anxiety of needles. The next two patients are the unconscious counterparts of these two: the Unconscious Christian Scientist and the Unconscious Anxious Patient both do not give consent, because they are unconscious, nor would either of them have given consent had he or she been conscious (but for different reasons, as above). And compare these, for control, to the person I'll just call the Unconscious Patient—who is unconscious but is neither a Christian Scientist nor anxious of needles, and who would have happily consented to the transfusion had she been conscious.

On entirely intuitive grounds, I submit that it's morally permissible (perhaps even required) to administer the transfusion in the case of the Unconscious Patient and also in the case of the Unconscious Anxious Patient. I am not sure about the other cases. But I'm quite sure that there's a morally relevant difference between the Unconscious Christian Scientist and the other unconscious patients, as well as between the Conscious Christian Scientist and the Conscious Anxious Patient. It may be permissible to administer the transfusion in the case of the Conscious Anxious Patient but not in the case of the Conscious Christian Scientist—but I'm not sure. Even if it's impermissible in both—or even permissible in both—there's still an important difference. It's at least more morally problematic to administer the transfusion to the Conscious Christian Scientist than to the Conscious Anxious Patient. And I would say something similar about the Unconscious Christian Scientist—it's at least more problematic to administer the transfusion to her than it is to the Unconscious Anxious Patient.

What explanations—hopefully, vindicating explanations—can we offer of these intuitive judgments?

Talk of whether the unconscious patients would have consented had they been conscious may explain the difference between the Unconscious Patient and the two other unconscious ones (the Anxious Patient and the Christian Scientist), but not between those two—both of them would not have consented had they been conscious. So this won't help. Nor will talk of what they would have consented to had they been rational, for we can stipulate that the (conscious or unconscious) person who is anxious of needles would not have been so anxious had he been rational, and also that the (conscious or unconscious) Christian Scientist would not have been a Christian Scientist had she been rational. Even under these assumptions, an intuitive difference between them remains. So idealizing on rationality won't do the work needed here. Nor will it help to talk about how the different patients will feel about things retroactively. This is a tempting line of thought, because we can imagine an Unconscious Anxious Patient who—when he regains consciousness—is happy to find out that he received the transfusion while unconscious (we can even easily imagine a Conscious Anxious Patient who is retroactively happy about having received the transfusion against his will at the time), and we can easily imagine an Unconscious Christian Scientist who—when she discovers that she received the blood transfusion when unconscious—is deeply troubled by this. But we shouldn't let this mislead us. Even if the religious commitments of the Christian Scientist are such that she can be (and indeed is) perfectly happy retroactively with the transfusion,³⁷ the intuitive difference between the two Christian Scientist patients on one hand and the two anxious patients on the other survives.³⁸ So we need another explanation.

A natural suggestion is that there is an important difference between the role that the commitment to Christian Science plays in the life of the Christian Scientist and the role that the needle anxiety plays in the life of those anxious of needles. The needle-anxious person may think of his anxiety as something external to his self, as something he finds himself with, rather than something he does, or something he is; he may feel alienated from it, or he may think of it as something to overcome.

37. Perhaps Christian Science only offers a kind of deontological constraint against receiving the transfusion, one that doesn't say that the state of affairs of having received one is at all bad.

38. If you're not convinced, consider the following: Suppose our Unconscious Anxious Patient will be terribly disturbed even to find out retroactively that he was given the transfusion. Still, there's an important difference between him and the Christian Scientist (who would also be retroactively upset), and one way of seeing this is thinking about deceiving the two about the treatment they received. I think it's much less morally problematic to hide the fact of the transfusion from the Anxious Patient than from the Christian Scientist. What explains this difference? I think that here too the point about depth that I am about to get to in the text does the work.

The Christian Scientist—or anyway, the one we’re talking about now—is not like this. She is committed to Christian Science; she identifies with it, she endorses this commitment; this commitment is a part of her self-conception, a part of who she is, of how she thinks of herself. The Christian Scientist’s commitment to Christian Science lies very close to the core of her self, whereas the needle anxiety is much more peripheral for the other patient.

One way of making such metaphors more precise is by employing higher-order desires and attitudes.³⁹ The Christian Scientist wants not to receive the blood transfusion. She also has, we can safely assume, higher-order desires that endorse this desire—she wants to continue having that desire, she wants that desire to bring her to action, and so on. The typical person who is anxious of needles will be very different in these respects—he will want not to receive the transfusion, but he may also want to rid himself of this first-order desire, he may want that desire not to bring him to action (he may want, say, his *joie de vivre* to outweigh his anxiety), and so on. Thus, while both patients desire—perhaps equally and maximally intensely—not to receive the blood transfusion, this desire is superficial in the case of the needle-anxious patient, whereas it is deep and central to the self-conception of the Christian Scientist. The Christian Scientist’s desire not to receive the blood transfusion is a part of a coherent, harmonious system of desires and wholehearted commitments, and this just isn’t the case with the needle-anxious patient.

This means that administering the blood transfusion in the case of the (conscious or unconscious) Christian Scientist amounts to a kind of an assault on the self in a way that a similar action in the case of the

39. Thinking about higher-order desires as a way of making sense of thoughts of endorsement—and indeed the self—is, of course, central to the work of Harry Frankfurt. See, e.g., Harry G. Frankfurt, “Freedom of the Will and the Concept of a Person,” *Journal of Philosophy* 68 (1971), reprinted in Harry G. Frankfurt, *The Importance of What We Care About* (Cambridge: Cambridge University Press, 1988), 11–25. And for such thoughts in the more specific context of discussions of autonomy, see generally Gerald Dworkin, *The Theory and Practice of Autonomy* (Cambridge: Cambridge University Press, 1988).

As one *Ethics* editor noted, in response to objections (mostly, from Watson) Frankfurt had to revise his account and include—as a way of stopping the ever-higher-order regress—something like a decision or a decisive commitment (see Harry G. Frankfurt, “Identification and Wholeheartedness,” in *Importance of What We Care About*, 159–76). It’s not clear that such a move can achieve what Frankfurt needs it to achieve, nor is it clear how it would apply to cases like that of the Unconscious Patient (whether it does will depend on whether we’re talking about a general, standing kind of decisive commitment, which may be present there, or a decision at that point in time, which is not). But these further issues need not concern us here: whether or not Frankfurt is right about the centrality of higher-order attitudes (perhaps coupled with a decisive commitment) in understanding freedom and responsibility, he’s clearly right about harmony across the hierarchy of attitudes mattering in something like the way utilized in the text—that is, as capturing the intuitive idea of a depth of a commitment.

needle-anxious patient does not.⁴⁰ In both cases, administering the transfusion amounts to a frustration of an expressed desire (in the cases of the conscious patients) or a hypothetical one (in the case of the unconscious ones). In both cases, if you administer the transfusion, you will be treating the patients in a way that usually requires consent, without their consent (and perhaps in the face of a withholding of consent). This is already problematic. But in the case of the Christian Scientist you will also be launching an assault on her very self. It is no surprise that this seems more seriously problematic, then.

This also nicely explains why the patient's being unconscious seems to solve the problem entirely in the case of the Anxious Patient, but not in the case of the Christian Scientist. With the Anxious Patient, the problem is a superficial one, and superficial solutions suffice to deal with it. And his being unconscious is a step in the right direction—it doesn't do away with his desire (not to receive the transfusion), but it removes the immediate objection and the need to, well, hold him down. With the Christian Scientist, the problem is deep, and so superficial solutions are just not good enough. Even when she's unconscious, treating her in a way that goes against her deep commitment, in ways she would have rejected had she been conscious, still amounts to an assault on her self. To put it a little dramatically: If the Conscious Anxious Patient suddenly becomes unconscious, the main problem in administering the transfusion (thereby saving his life) has just been solved. If the Conscious Christian Scientist becomes unconscious, at least one main problem is still very much in place.

Now, this explanation—in terms of depth, alienation, centrality to the self, and higher-order desires—goes a long way toward the vindicating explanation we were looking for of the intuitive judgments we started with. Still, let me concede its limitations. First, it does not cover all cases. Think, for instance, of the normative relevance of consent to sex. This seems to be the paradigmatic case where we require (at least) actual consent. At least in the face of actual refusal, no hypothetical consent will ever do—there is no way of completing the sentence “True, she said ‘no,’ but she would have consented had . . .” in a way that renders sex permissible.⁴¹ This is so even if her refusal on the specific occasion was based

40. At various points, Frankfurt speaks of the decisive commitment to identify with a desire as a way of constituting the self. At one point (Frankfurt, “Identification and Wholeheartedness,” 171) he talks of such identification as taking responsibility for one's characteristics, thereby becoming responsible for one's character. He doesn't develop this point. For a related discussion of taking responsibility as a voluntary action that renders one responsible, see David Enoch, “Being Responsible, Taking Responsibility, and Penumbral Agency,” in *Luck, Value and Commitment: Themes from the Ethics of Bernard Williams*, ed. Ulrike Heuer and Gerald Lang (Oxford: Oxford University Press, 2012), 95–132.

41. Estlund, *Democratic Authority*, 216. Even here, though, there are complications. Perhaps there are cases where hypothetical consent (not in the face of actual refusal)

on unendorsed desires, or some such. I return to this example in the next section. Second, I have not ruled out all possible competing explanations of even just the five patient cases I've been discussing. I've ruled out some, and I've done what I can to make the one in terms of depth plausible. But I should note that other explanations are possible, and that if one is put forward, it should be evaluated on its (comparative) merits.

Even with these restrictions in mind, then, some progress has been made. For we seem to have made a plausible case for the following claims: At least sometimes, hypothetical consent does matter, because treating people according to the way they want or would want to be treated is important, at least when their so wanting to be (or not to be) treated is an endorsed desire, one they identify with. And in those cases, it's important to treat them in such a way because treating them differently can amount to an assault on their self.

We can say a little more now. Respecting people's desires and commitments is sometimes a way of responding appropriately to their autonomy—to the fact that, at least often, being the author of one's own life story is an aspect of living a good life, and perhaps also that people have a right to be—to an extent—the authors of their life story.⁴² The reason we have to not administer the blood transfusion to the Conscious Christian Scientist is precisely the reason we have to respect her autonomy. (We also have reasons to save her life. It's not obvious to me which is the weightier reason, and it may vary with context. What is clear, though, is that we have some reason not to administer the transfusion, and that that reason has to do with the patient's autonomy.) The value of autonomy is not indifferent to the distinction between the Christian Scientist and the Anxious Patient: When you administer the transfusion to the Conscious Anxious Patient against his expressed desire, you are most certainly treating him in a way that he doesn't consent to. But because he is alienated from his anxiety, because he doesn't identify with it, because it is not central to his self, you are not harming his ability to write his own life story in anything like the way you do this if you give the Conscious Christian Scientist the transfusion. And the value of autonomy also explains why hypothetical consent sometimes matters. For sometimes the way in which you respect someone's self-authorship is by treating them

can matter—think here, e.g., of waking up a romantic partner, in the context of a long relationship, by touching them in a sexual way that would require actual consent in other settings. Perhaps here it's permissible to do so because the partner would have consented had he or she been awake. And even if hypothetical consent can never make sex permissible, perhaps it can still make a difference, say, in how seriously impermissible it is. Still, the point in the text by and large stands.

42. It is hard to explain the value of autonomy in nonmetaphorical ways. For this metaphor, see Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986), chap. 14.

in ways they would consent to, under some hypothetical conditions—that is, under hypothetical conditions that will bring out and emphasize the desires and attitudes with which they identify, and that will perhaps filter out, to an extent, those from which they are alienated.⁴³

VII. NO NORMATIVE WORK, AGAIN?

You may still be worried, though, that on the emerging picture hypothetical consent doesn't do any normative work. You may think, perhaps, that what does the normative work in such cases is not whether or not the relevant person would have consented, but rather what their deeper commitments entail. And so, perhaps, hypothetical consent again drops out of the picture. Let me make, then, the following points in response.

First, treating people in accordance with their deeper commitments is very close to treating them in the ways they would consent to, under suitably characterized hypothetical conditions.⁴⁴ Notice that the value of autonomy works precisely in the way described above, in Section IV, as the way of avoiding the objection posed by the Transitivity Argument: autonomy dictates that the patient's deep commitments—whatever their precise details—be respected. This means that we have a kind of a multiple realizability going on. What does the normative work is not the patient's specific commitment to Christian Science, as other deep commitments that would lead her to refuse treatment (being a Jehovah Witness, say) would work just the same. The only thing that is in common to all the different commitments by the patient that would have a similar normative effect is precisely that they are her deep commitments, that they are what she would act on (in the suitably characterized hypothetical conditions). Notice that the in-virtue-of chain is broken also by another point noted in Section IV, namely, the reasons underlying the hypothetical consent (or lack thereof) are reasons for the patient,⁴⁵ whereas it's the value of autonomy and the hypothetical consent that are the reasons for others—the reason for which the doctor may avoid administering the

43. I want to emphasize that we are not doing politics here—not yet, that is. So of course there are serious political problems in trying to institutionalize the points in the text—perhaps, for instance, we do not want the state to be in the business of diagnosing citizens' deep commitments, or of distinguishing between central and peripheral parts of its citizens' selves. But the seriousness of such concerns does not challenge the points made in the text—at a much more abstract stage of the discussion, as it were.

44. If commitments are understood at least partly in dispositional terms (as seems plausible), and if dispositions are understood at least partly in counterfactual terms (as seems almost inevitable), then the relation between the deep-commitment story and the hypothetical-consent story may be tighter still.

45. And again, they may not even be good, normative reasons; they may just be motivating reasons, causally but not normatively grounding the patient's attitudes.

blood transfusion is not Christian Science doctrine, of course. (Presumably, the doctor doesn't take that to be a [normative] reason at all, and rightly so.) Rather, it's that the patient would not have consented had she been conscious, and that respecting this hypothetical refusal is a way of respecting the patient's autonomy.⁴⁶ In this way, then, the deep-commitments story and the hypothetical-consent story are both deeply personal—they are not about, say, which deep commitments are genuinely reason giving, but about their being taken to be reason giving by those whose commitments they are. So it's much more natural to understand the deep-commitment story not as a competitor of the hypothetical-consent story but rather as an explanation thereof. And if you are still inclined to view these stories as competitors, I am tempted to say that I no longer care which one wins. The supporter of the normative role for hypothetical consent can still get all or most of what she wants from the deep-commitment story, suitably understood.

Second, we may want to try to force a wedge between these two stories. A nice way of doing this would be to consider a Weak-Willed Christian Scientist.⁴⁷ This is someone who is as seriously committed to Christian Science as anyone but who, in the face of a substantial risk of death, gives in to temptation to go on living and asks for that blood transfusion. We can consider different cases. First, there's the Conscious Weak-Willed Christian Scientist, who consents to the transfusion, against his deep commitments. In such a case it seems to me clear that we should administer the transfusion, but that this is so may be the result of factors that are irrelevant for our discussion (perhaps that Christian Science doctrine is wrong, that the reasons for him to go on living are very good ones, or perhaps because of the value of sovereignty, to which I turn below). Then there's the Uncon-

46. For this reason we can also see that the normative relevance of hypothetical consent (or lack thereof) here is, in Waldron's terms from Sec. II, not just a matter of reason; it's also about will. But it's about will in a way that leaves room for the relevance of hypothetical consent.

47. I owe this example to Alec Walen. There are other cases in which the relation of hypothetical consent to deep commitments may not be as clear. Think of a case, for instance, in which the patient, in some sense, deeply wants to receive the transfusion (because she wants to survive), but she would never give her consent, because she believes that consenting to the transfusion—but not necessarily receiving it—is what's forbidden by her religion. (I thank Ram Rivlin and an editor for *Ethics* for this kind of example.) I'm not sure what to say about such a case—it's clear to me neither what this patient is most deeply committed to nor whether she should be given the transfusion (either when she's conscious or when she's unconscious). Perhaps for our purposes here we can just restrict the discussion to cases in which the reasons for consent to something have to do with that something rather than with the consent itself. (These are standard cases in a sense analogous to that of cases where one's reasons for intending to φ are very close to one's reasons for φ -ing, rather than the nonstandard, toxin-puzzle cases, where the reasons for intentions are divorced from the reasons for the action intended.)

scious Weak-Willed Christian Scientist—someone who is deeply committed to Christian Science but who would have, had he been conscious, given in to temptation and consented to the administration of the life-saving transfusion. This is where the deep-commitment story generates a result that differs from that generated by the hypothetical-consent story: following his deep commitments would dictate not administering the transfusion, or at least treating this case as no different from that of the strong-willed Christian Scientist; following hypothetical consent would dictate administering the transfusion, and perhaps ignoring entirely the deep commitment to Christian Science (as we presumably would in the case of the Conscious Weak-Willed Christian Scientist). In this kind of case, then, which verdict is more intuitively plausible? I have to admit that this doesn't seem clear to me. Perhaps the most intuitively plausible thing to say is that the case of the Unconscious Weak-Willed Christian Scientist differs in important ways from both that of the strong-willed Christian Scientist (who wouldn't have given consent) and that of the non-Christian Scientist (who would, without any conflict with his deep commitments).⁴⁸ If so, both the deep-commitment story and the hypothetical-consent story do normative work here. And if so, it's not the case that all the work is done by the deep-commitment story. Hypothetical consent is not epiphenomenal after all.

VIII. AUTONOMY AND ACTUAL CONSENT: SOVEREIGNTY AND NONALIENATION

A picture begins to emerge, then.⁴⁹ Because of the value of autonomy, hypothetical consent sometimes matters. Sometimes, treating someone

48. The case of the Unconscious Weak-Willed Christian Scientist is a particular instance of a more general schema, where the deep commitment is there but is somehow not manifested in decision in the possible world on which (hypothetical) consent is sought. Sarah Moss suggested to me another particular instance—in her case, the Christian Scientist, had he been conscious, would have (for some reason) lost his faith. It's not clear to me whether in this case our intuitions are similar to the ones regarding the case in the text. Nor is it clear to me that this is a case we need to worry about in our context; it may be a case of "masking a disposition" (see Mark Johnston, "How to Speak of the Colors," *Philosophical Studies* 68 [1992]: 221–63), which calls for a general treatment and doesn't raise problems that are peculiar to our context.

For related reasons, we may also want to employ a move common elsewhere and switch from asking (e.g., about a patient) what she would do in the relevant hypothetical conditions to asking what her ideal advisor would advise that be done to her, nonideal as she is. This may be needed in order to deal with cases like Pallikkathayil's example of the absent colleague who, if asked, would gladly consent to your borrowing his book, but who would resent your borrowing the book if no (actual) consent were given. See Japa Pallikkathayil, "Hypothetical Consent Reconsidered" (unpublished manuscript).

49. Note that the discussion of autonomy that starts in the text here remains entirely in the normative-ethics domain. I do not engage here the discussion that ties this value to

in a way that he doesn't consent to but that he would consent to under hypothetical conditions that filter out his superficial desires and emphasize his deep commitments does not amount to an attack on his self or to an affront to his ability to be the author of his life story—anyway, not as much so as similar treatment of someone who wouldn't even consent under those conditions.

Recall now the general point about idealization from Section V. To show that hypothetical consent ever has the kind of normative significance that actual consent has, we need to see why actual consent matters and then see whether hypothetical consent can also satisfy these considerations. And partly, this is what we've been doing in the previous section. For surely at least one reason for which actual consent sometimes matters is because of the value of autonomy. And it is the upshot of the previous section that sometimes the value of autonomy grounds the significance of hypothetical consent just as it does that of actual consent.

So far, though, I've been focusing on only one way in which autonomy may matter. According to this way, I've been insisting, a violation of your autonomy amounts to a kind of alienation—you are not allowed to write your own life story according to your deep commitments. But autonomy may matter in another way as well.⁵⁰

Suppose that my daughter, out of concern for my health, will sometimes hold on to the salt at the dinner table and refuse to pass it along, even when I ask her. I appreciate the gesture and the sincere concern, of course. But sometimes, I insist. I can explain that I understand the health issues, but that it's my body, and my life, and that I want the salt, please. I can, it is natural to say, assert my autonomy. At least in some such cases, if I so insist and my daughter still refuses to let me have the salt, she is offending against my autonomy. But notice, of course, that nothing like alienation is at all involved. My deep commitments are in no way threatened by my daughter's refusal to pass the salt, nor is my ability to write my

the metaphysics of agency—I don't think I have to. For an example of this discussion, and for many references, see Sarah Buss, "Personal Autonomy," in *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Stanford, CA: Stanford University, 2002), <http://plato.stanford.edu/entries/personal-autonomy/>.

50. Of course, nothing depends on the word "autonomy." It doesn't matter whether my distinction here is between two ways in which autonomy matters (two distinct values or goods, e.g., or perhaps two distinct deontological rights, or some such) or between autonomy and another, closely related value.

Kuflik ("Hypothetical Consent," 147–49) also stresses the relation between hypothetical consent and autonomy. But he doesn't distinguish between the two values I am about to distinguish between, and perhaps as a result he doesn't even address the obvious difficulty: seeing that no actual decision has been made by the relevant agent, how does respecting a decision she should have made amount to respecting her authority?

I don't know of anyone drawing the exact distinction I draw in the text. There is a similar distinction (in other terms) in Pallikkathayil, "Hypothetical Consent Reconsidered."

life story in accordance with them. The way in which my autonomy is violated is different: at least after having asserted my autonomy, the salt in my dinner is something over which I should have the last say; it's within my area of sovereignty, so to speak. This too is a way in which autonomy is sometimes important, and this too is a way in which it should sometimes be respected.

I'm going to distinguish, then, between autonomy reflecting a concern for nonalienation and autonomy reflecting a concern for sovereignty. Though the two are closely related, they are still distinct. I hope to say more about the relation between these two in future work. Perhaps, for instance, one of the two is more basic, and the other derivative. Perhaps, for instance, nonalienation is the fundamental thing that is of value here, and sovereignty is significant only because usually, in general, perhaps in a specific cultural setting, it goes hand in hand with nonalienation. Or perhaps the two are both, independently, important. These are interesting questions. But for our purposes here we don't need to answer them. For our purposes here it's sufficient to note that however nonalienation and sovereignty are related, and whatever exactly their nature (as values narrowly understood, or as normatively relevant features more generally, perhaps with a deontological aroma), they are distinct; that autonomy is associated with both; and that, therefore, actual consent matters sometimes because of nonalienation, sometimes because of sovereignty, and, presumably, sometimes because of both.

Notice also that these two kinds of concern—nonalienation and sovereignty—are very closely related to Waldron's distinction between reason considerations and will considerations. But nonalienation need not go hand in hand with reason. Even if reason rejects Christian Science without reservation, still there is reason to respect the Christian Scientist's autonomy, at least within certain restrictions.

We can now return to hypothetical consent. If there are two distinct reasons why actual consent matters when it does—nonalienation and sovereignty—we can now ask whether hypothetical consent can answer to these two concerns as well as actual consent presumably can (at least sometimes). And here, the distinction between the two ways is important. For the answer seems to me to be "yes" in one and "no" in the other.

When it comes to nonalienation, as we've seen in the previous section, sometimes hypothetical consent can do the work. So when actual consent is sought because of concern with alienation, hypothetical consent may be a nonpale substitute.⁵¹

51. Or perhaps we shouldn't think of it as substituting for actual consent at all. For this claim, but in a very Kantian context, see Hill, "Hypothetical Consent in Kantian Constructivism," 320–22.

When it comes to sovereignty, though, this is not so. If I insist on the salt decision being mine to make, there's nothing by way of hypothetical consent that can respond to this. If the decision is mine to make and I've actually decided one way, the fact that I would have decided it differently under different conditions is just neither here nor there. If my daughter should respect my salt decision, she should respect the decision I actually make. When it comes to sovereignty, then, hypothetical consent is not even a pale form of consent. It really is no consent at all.

The distinction between the two autonomy concerns—that of non-alienation and that of sovereignty—nicely explains the initial state we find ourselves in regarding hypothetical consent. As I emphasized at the outset, we seem to be torn here—intuitively, hypothetical consent does seem to make a difference in many cases, and yet when we think about it, it's not clear how it can be more effective normatively than hypothetical water is in quenching thirst. Perhaps the initial tension is due, then, to our failure to distinguish between the nonalienation and sovereignty concerns, both relevant to consent. The sovereignty concern is the one that gives rise to the thought that hypothetical consent cannot make a difference. And the normative relevance of nonalienation explains why in many cases we do seem to care about hypothetical consent.

I don't want to suggest that sovereignty is easy (and as I said, I hope to have more to say about it in future work). There may be cases in which sovereignty too may amount to more than just whose say decides what. Consider a composer who seems to have made a mistake in the score of the music that she wrote. There's a case to be made that by correcting it, we are in effect respecting her sovereignty, not (or not just) her deep commitments.⁵² Sovereignty too may require some (gentle) idealization, perhaps even relying on some hypothetical attitudes of the relevant person (though probably not on her hypothetical consent). Still, it's clear that sovereignty is very different from nonalienation. If the composer is standing there, insisting that it's not a mistake, we can try to convince her that it is; we may want to change the score, her protests to the contrary notwithstanding; we may even insist on doing this in order to respect her deep commitments (to which she is somehow not being loyal now). What we cannot plausibly do, though, is insist that we are thereby respecting her sovereignty.

Let's revisit our examples. With regard to the autonomy of conscious patients, we seem to be concerned both with nonalienation and with sovereignty. So the reason it's more problematic to administer the transfusion to the Conscious Anxious Patient than to a conscious patient who fully consents seems to have to do with the concern for sovereignty. And the reason it's even more problematic to administer the transfusion to

52. I owe this example and the point that follows to Peter Railton.

the Conscious Christian Scientist is that in her case we're not just concerned with sovereignty, but also with nonalienation. As we shift to the unconscious patients, talk of sovereignty seems no longer relevant in the same way. So there's no serious problem with the Unconscious Patient or, interestingly, with the Unconscious Anxious Patient. The nonalienation element remains, however, in the case of the Unconscious Christian Scientist, and so in her case the situation is still morally problematic (and what we should do, all things considered, may depend on the specifics of the case).

How about the example of consent to sex, where it seems like a fairly robust intuition that no merely hypothetical consent can ever make sex morally permissible? In that case, it seems that the interest we have in sovereignty is of tremendous importance. When it comes to sex, it's important that each person has the final say on her or his participation in sexual practices. Though we may also care about nonalienation when it comes to sex—perhaps this is one of the reasons we often think that superficial actual consent is also not enough—we care mostly, and very strongly, about sovereignty.⁵³ And when this is what consent is supposed to do, we already know that hypothetical consent is no substitute at all.⁵⁴ This example, then, also confirms the account above.

Finally, consider the political case. When we care about liberty and autonomy in the political sphere—the kind of thing we care about when we worry about the tension between authority and liberty, or about people treating each other as free and equal, or about nonsubordination—do we care about nonalienation, or about sovereignty, or about both? Well, it seems to me we care about both. But I am especially confident

53. You may think that we care about sovereignty whenever bodily integrity is involved, not just when it comes to sex. And this may blur the difference between the sex case and the blood transfusion cases (where bodily integrity is also at issue). But there are many differences. First, not all cases of bodily integrity are equally important in terms of sovereignty. Sex cases seem special in this regard. Second, the reason to administer the transfusion is powerful and benevolent—for the good of the person whose bodily integrity is in question. This is not the case in consent-less sex cases. Third, with blood transfusion, consent seems all we want. Perhaps we are even okay with mere absence of objection (in suitable conditions). When it comes to sex, we may require more—like a positive attitude, and not merely the absence of a negative one, or the expression of a neutral one. This is why with some transfusion cases unconsciousness seems to make things easier, not harder. Not so with sex, of course.

54. Think about it in terms of an ideal-advisor model again. Here, we can ask whether one's ideal advisor would want or consent to sex (on behalf of one's actual, nonideal self). But here it becomes crucial to ask also whether one's ideal advisor would want one to have sex given actual nonconsent (of the actual, nonideal self). And it seems very plausible to answer this question in the negative. This is another way of showing, I think, the central role for sovereignty here. Some of our deepest commitments about sexual contact have to do with the significance of sovereignty. It's not at all clear that we would get a similar result for the ideal-advisor thought experiment for other infringements of bodily integrity.

we care about sovereignty. At least partly, this is about who gets to decide (e.g., over me). The thought, for instance, that in a real-life political disagreement that results in an impasse you can coerce me, insisting that there's really no problem with my liberty (or autonomy or some such) because the policy you're pushing down my throat is one that is approved of by my own deeper commitments, or one that I would accept under suitably described hypothetical conditions—that thought just seems entirely ludicrous.⁵⁵ Perhaps it's not quite as ludicrous as the thought that hypothetical consent suffices for the permissibility of sex, but it's close.⁵⁶ I am willing to concede that if a political arrangement was in line with everyone's deep commitments, that would have counted in favor of it; but still, a dictatorship that governs in a way that jibes well with all the governed's deep commitments but in which the governed have no say is still very much a dictatorship, and not all there is well in autonomy terms. So sovereignty too is of utmost importance here. If this is so—if the autonomy relevant to political authority is centrally about sovereignty—and if the discussion above gets it right, then what follows is that hypothetical-consent accounts of legitimate political authority are hopeless.⁵⁷ I think of this result as an advantage of my account of the moral significance of hypothetical consent.

IX. WHEN DOES HYPOTHETICAL CONSENT MATTER?

Without pretending to say something conclusive, or complete, or very general, I think we can conclude with the following lessons: First, even if hypothetical consent doesn't often do the work actual consent does, still it can be normatively significant in other ways. Second, and even restricting ourselves just to substituting for actual consent, a conclusive argument against the normative significance of hypothetical consent remains to be made. The two arguments most common in the literature raise important challenges but are far from being conclusive. The way out of the Transitivity Argument is the multiple-realizability line, and the way out of the thought that hypothetical consent is no substitute for actual consent is to note that whether or not it is depends on why it is that actual consent matters, when

55. This is a point I make much of in David Enoch, "The Disorder of Public Reason: A Critical Study of Gerald Gaus's *The Order of Public Reason*," *Ethics* 124 (2013): 141–76; see also Enoch, "Against Public Reason."

56. I deliberately chose a nonpolitical example (the salt case) in order to show that sovereignty autonomy is relevant not just in politics. The point in the text is that though it's not relevant only in politics, it most certainly is relevant in politics.

57. Dworkin's influential quote—that a hypothetical contract is not a pale form of a contract but rather not a contract at all—is not about consent exactly, but about contractual obligation. Still, it may be interesting to point out that in contracts, too, the relevant value of autonomy seems to be much more about sovereignty than about nonalienation. Perhaps this is why Dworkin's dismissal of hypothetical contracts is so plausible.

it does. Third, actual consent sometimes matters because of nonalienation, and sometimes because of sovereignty. Hypothetical consent can (sometimes) suffice for nonalienation, but not for sovereignty. Fourth, one natural way of understanding nonalienation is in terms of Frankfurtean endorsement, itself understood in terms of higher-order attitudes.

This is not quite a recipe for determining when hypothetical consent matters and how. But it's a good start, I think. For we now know—if only roughly—what you must show in order to show that in a specific case hypothetical consent can be normatively significant. And this too, it seems to me, may be significant progress.