Why Decision-making Capacity Matters

[*penultimate draft; please cite published version: https://doi.org/10.1163/17455243-20213610*]

Ben Schwan (corresponding author)
ben.schwan@case.edu
Department of Bioethics
Case Western Reserve University
Cleveland, OH, USA

Abstract

Decision-making Capacity (DMC) matters to whether a patient's decision should determine her treatment. But why it matters in this way isn't clear. The standard story is that DMC matters because autonomy matters. And this is thought to justify DMC as a *gatekeeper* for autonomy—whereby autonomy concerns arise if but only if a patient has DMC. But appeals to autonomy invoke two distinct concerns: concern for *authenticity*—concern that a choice is consistent with an individual's commitments; and concern for *sovereignty*—concern that an individual exercises control over that which is hers to control. Here, I argue, neither concern can alone explain why DMC matters. Instead, DMC matters because it indicates a *harmony* between the two concerns—the demands of each concern are more likely to agree if a patient has DMC. This vindicates the standard story, but also makes clear that DMC is an inappropriate gatekeeper for autonomy.

Keywords: decision-making capacity, autonomy, authenticity, sovereignty

1. Introduction

Consider the following cases:

> **Jill**. Jill has cancer. She's exhausted all her treatment options but for these: i) a round of chemotherapy that could prolong her life by up to two years; ii) palliative care while the cancer runs its course—giving her another six months or so. Jill understands her options and appreciates her situation; she weighs the pros and cons and opts for the chemo.
>
> **Jack**. Jack has cancer too. He's exhausted all his treatment options but for these: i) a round of chemotherapy that could prolong his life by up to two years; ii) palliative care while the cancer runs its course—giving him another six months or so. In addition, due to complications from a broken crown, Jack is cognitively impaired such that he neither understands the difference in time between his two options, nor appreciates the suffering he'd experience during the chemo. He too opts for the chemo.

Intuitively, we should do as Jill wishes. I'm not sure what we should do with Jack. We'd need more information, I think. I'm at least confident about this, though: Jack's decision does not *determine* what we should do in the way that Jill's does. And the explanation for this has something to do with the abilities that Jill has and Jack lacks—Jill is able to understand the relevant information, appreciate how it applies to her, reason about her options, and communicate her choice. Jill, in bioethics-speak, has *decision-making capacity*, and Jack does not. And these facts about Jack and Jill help explain why Jill's decision should determine her treatment while Jack's should not.

So, decision-making capacity matters. More specifically, it matters to whether a patient's decision should determine her treatment. Or, as I'll sometimes put it, decision-making capacity matters to whether a patient's decision should have *authority*.[1] But *exactly why* it matters in this way is less clear. The standard story appeals to autonomy—decision-making capacity matters because autonomy

[1] To be precise, as I'll use the term, a person's decision has *authority* regarding a choice if and only if it determines (or will determine) which option prevails. Authority is, in this sense, descriptive—it's a claim about whether a person's decision *actually* determines (or will determine) which option prevails, not a claim about whether a person's decision *should* determine which options prevails.

matters. And this standard story is thought to ground the use of decision-making capacity as a *gatekeeper* for autonomy, whereby a patient's decision generates autonomy concerns if but only if that patient has decision-making capacity. But, as Brudney and Lantos (2011) and Enoch (2017) have recently argued, appeals to autonomy invoke two distinct concerns: concern that an individual's life goes in a way that is consistent with her deep commitments and concern that an individual exercises control over that which is properly hers to control. And, in light of these distinct concerns, the standard story about why decision-making capacity matters can be told in multiple ways.

Here, I show that neither concern can alone explain why decision-making capacity matters. Instead, I argue, decision-making capacity matters because it indicates a *harmony* between the two autonomy concerns—since a patient with decision-making capacity is likely to decide in accordance with her deep commitments, to ensure that she exercises control over a choice is also (at least likely) to ensure that the chosen option conforms to her deep commitments. This vindicates the standard story. But it also suggests that decision-making capacity is an inappropriate gatekeeper for autonomy—by deferring to the decisions of all and only patients who have decision-making capacity, we risk ignoring important tradeoffs between the two autonomy concerns, both in those with decision-making capacity and in those without it.

I begin, in section one, by defining the bioethicist's notion of decision-making capacity. Then, in section two, I tell what I take to be the Standard Story about why decision-making capacity matters and explain a popular way of operationalizing this story: what I call the Gatekeeping Model. Next, in section three, I pull apart the two autonomy concerns and show how they render the Standard Story and Gatekeeping Model ambiguous. In section four, I explore three ways of resolving this ambiguity and defend my view that decision-making capacity matters because it indicates a harmony between the two autonomy concerns. Finally, in section five, I suggest that in light of my reading of the Standard Story, we should abandon the use of decision-making capacity as a gatekeeper for autonomy.

2. Decision-making Capacity Matters

Decision-making capacity, for present purposes, is best thought of as a term of art, employed by bioethicists, psychiatrists, clinicians, and researchers to pick out a very specific set of abilities:

> **Decision-making Capacity** (DMC): A person has **decision-making capacity** regarding a particular set of options if and only if she has the ability to
>
> - **understand** the relevant information,
> - **appreciate** her situation,
> - **reason** about her options,
> - and **communicate** her choice.[2]

There's some debate about how to understand each of these conditions and about exactly where the relevant thresholds lie.[3] But for present purposes, a rough characterization will do. To understand the relevant information is (roughly) to know the facts about the options that are relevant to what you care about. To appreciate your situation is (roughly) to recognize your condition and that your decision and its consequences apply to you. To reason about your options is (roughly) to engage in rational deliberation about the pros and cons of each. And to communicate your choice is (roughly) to behave in some way that reliably indicates your preferred option to others.[4]

Though rough, it's important for present purposes that this characterization of DMC is tightly tethered to how DMC is assessed and measured in medical research and clinical care. Researchers, clinicians, nurses, psychiatrists, and ethicists are frequently tasked with determining whether a patient's decision should have authority regarding some choice *based on* an assessment of that patient's abilities to understand, appreciate, reason, and communicate. These assessments of DMC are sometimes formal—using, for example, MacArthur Instruments or the Capacity to Consent to Treatment

---

[2] For a more thorough account, see Appelbaum (2007).

[3] In fact, it's typically thought that the thresholds *change* from case to case depending on what's at stake. For a helpful overview of the relevant literature, see Kim (2009, Chapter 6).

[4] One additional feature worth noting: as the "regarding a particular set of options" clause suggests, DMC is *local,* not *global*—a person has (or lacks) DMC *with respect to a particular decision or set of options*. So, for instance, one might *have* DMC regarding which medication to take, while *lacking* DMC regarding whether to have surgery. But this feature of DMC won't play a role in what follows.

Instrument[5]—and sometimes informal. But in any case, they have certain predictable results. For instance, mental health conditions, such as depression and schizophrenia, and physical illness and pain often undermine a patient's DMC.[6] A recent analysis by Sessums, Zembrzuska, and Jackson (2011) determined that 2.8% of healthy elderly adults, 20% of patients with mild cognitive impairment, 26% of medicine inpatients, and 68% of learning disabled patients lack DMC.[7] As mentioned above, there's no general agreement about how to understand the abilities that constitute DMC or about where the relevant thresholds for each ability lie. But note that to get outcomes like these, the thresholds can't be set too low (because, otherwise, far more people would be capacitated) or too high (because, otherwise, far fewer would be). So while I'll remain noncommittal about the precise nature of the abilities and their thresholds, I will have in mind a generic sense of DMC that reflects how it's used in practice such that, for example, 2.8% of healthy elderly adults lack it.

In addition to this tether to practice, it's also important to note that DMC, as I'm understanding it, is itself *non-normative*. To claim that someone has or lacks DMC, then, is to make a descriptive claim about her abilities, without thereby *presupposing* any implications for how she should be treated in light of those abilities. This is worth noting because DMC is sometimes given a *normative* gloss such that *what it means* for a person to have DMC is that she has each of the above abilities to such degrees that her decision *should* have authority.[8] To say that someone has DMC in this normative sense is to make two claims: 1) that she has each ability to some respective degree, and 2) that in virtue of the fact that she has these abilities to these degrees, her decision should determine her treatment. I mention this normative characterization only to set it aside. How and why facts about these abilities are relevant to

---

[5] For an overview of the various instruments used to measure DMC, see Vellinga et al. (2004).

[6] See Hatherley's (2019) discussion of Appelbaum and Grisso (1995), Grisso et al. (1995), Grisso and Appelbaum (1995), and Hindmarch, Hotopf, and Owen (2013).

[7] See also Kim (2009, Chapter 3), who surveys studies about rates of incapacity among hospital patients and finds them to range between 31% and 48%.

[8] For a helpful gloss of the normative/non-normative distinction, see Hawkins and Charland (2020, sec. 1.2). And for an example of someone who employs the normative characterization of DMC (and is refreshingly explicit about doing so), see Kim (2009, especially Chapter 6).

whether a person's decision should determine her treatment is exactly what's at issue in this project. Given this, I need some way to pick out the relevant set of abilities that doesn't bake in any commitments about the ways in which they matter to whether a patient's decision should have authority. Hence, in what follows, it's the *non*-normative sense of DMC that I'll have in mind.

Again, the foregoing sketch of DMC is rough in many respects. But, despite its roughness, it should be enough to make clear that a person who lacks DMC might nevertheless make a genuine decision. The reason is that decisions come (comparably) cheap. To decide just is to form an intention or to settle a practical question or to make up one's mind about what to do.[9] As such, it's a basic mental event that doesn't require the robust abilities of DMC—the love-struck groom, blinded by his affection and *unable to understand* his relationship's flaws, might nevertheless decide to go through with the marriage; the brash jock might decide to pursue a basketball career, despite her *inability to appreciate* the hardships that she'll endure; the paranoid anarchist, in the grips of a conspiracy theory, might decide to build his bunker, despite an *inability to reason* about whether the government is out to get him. More generally, it's not altogether uncommon for a person to make a decision that is uninformed (because she can't understand relevant facts), imprudent (because she can't appreciate her situation), irrational (because she can't reason well about the relevant issue), or inscrutable (because she can't communicate her choice).

Importantly, nothing I've said so far is a criticism of DMC—of how it's conceptualized or how it's used. The bundle of abilities that constitute DMC are perfectly intelligible, they're certainly important, and 'decision-making capacity' is a natural label for them. Rather, it's just to point out the unavailability of a very simple story about why DMC matters—namely, that it matters because without

---

[9] For these characterizations, see Mele (2017, Chapter 2), Hieronymi (2009), and Boyle (2011), respectively. The disagreements among them (and their competitors) aren't of interest here. The relevant point just is that on any plausible account of deciding, decisions are possible without DMC.

Also note that this gloss of what it is to decide is consistent with the possibility of some decisions being *non-autonomous*. Later, in §5.2, I'll argue that even *autonomous decisions* are possible without DMC. Here, though, my claim is more modest.

DMC a patient *literally cannot decide*. Decisions are possible without DMC. And, later in §4, I'll consider whether and in what way the decisions of incapacitated patients have moral weight. So it's important to be clear here at the outset that when I talk of 'decisions' and 'being able to decide', I'll always have the conceptually thinner notion in mind, and 'DMC' (and 'capacitated' and 'incapacitated') will always refer to the bioethicists-and-friends' more robust conception of the abilities described above.

Again, DMC, so characterized, matters. In at least some cases, the fact that a patient has DMC explains why her decision should determine her treatment, and the fact that a patient lacks DMC explains why her decision should not. And this is a plausible role for DMC to play. It explains our intuitions in cases like Jack's and Jill's, and it's widely endorsed by bioethicists, psychiatrists, researchers, and clinicians. It's also an *important* role for DMC to play. Often, how a patient should be treated in light of her decision is obvious. But, other times, it's not. And when it's not, our assessment of whether the patient has DMC can help guide our treatment of her. To do this work *well*, though, we need to understand exactly why DMC matters in the way that it does.

3. Why DMC Matters: The Standard Story and the Gatekeeping Model

The standard story appeals to autonomy.

The Standard Story: DMC matters because autonomy matters.

To be autonomous is to be *self-governing*—to have the capacity to act according to one's beliefs, values, aims, and intentions. And to act autonomously is to *exercise* this capacity—for one's commitments to combine in the appropriate way and bear the right relation to what one does.[10]

Autonomy is typically thought to be worthy of respect—imposing constraints on the ways that those who possess it may permissibly be treated. For example, in their canonical *Principles of Biomedical Ethics*, Beauchamp and Childress explain:

[10] This sketch of autonomy is meant to be ecumenical and is, as a result, is vague. But it should do for present purposes.

> To respect autonomous agents is to acknowledge their right to hold views, to make choices, and to take actions based on their values and beliefs. (Beauchamp & Childress, 2013, p. 106)

The fact that you're autonomous, then, changes the normative landscape—making certain ways of acting toward you permissible and others impermissible. Sometimes, these changes are complicated. But often they're simple. To preview an example from below, if you decide to have a hoagie for lunch, then (in the normal range of cases) it would be wrong for me to ensure that you have a salad. And the most natural explanation for why this would be wrong is that it would fail to respect your autonomy.

So, the Standard Story says that the reason DMC matters to whether a patient's decision should have authority is that (at least sometimes) respecting autonomy requires that we do as a capacitated patient decides. And this is a natural story to tell. Autonomy is a ubiquitous ethical concern—generally and especially in the medical context. But it's not *always* relevant. Sometimes autonomy concerns don't arise.

Here are Beauchamp and Childress just a little bit later:

> Obligations to respect autonomy do not extend to persons who cannot act in a sufficiently autonomous manner…because, for instance, they are immature, *incapacitated*, ignorant, coerced, or exploited. (Beauchamp & Childress, 2013, p. 108; emphasis added)

And, a little later still:

> *Competence or capacity judgments* in health care *serve a gatekeeping role* by distinguishing persons whose decisions should be solicited or accepted from persons whose decisions need not or should not be solicited or accepted. (Beauchamp & Childress, 2013, p. 111; emphasis added)

The claim here is that the Standard Story justifies the use of DMC as a *gatekeeper* for autonomy—determining when autonomy considerations come into play. When a patient has DMC, autonomy demands that her decision determines her treatment, but when a patient lacks it, autonomy makes no such demand because autonomy concerns simply don't arise. Call this reading of the Standard Story the Gatekeeping Model.

> **The Gatekeeping Model**: A patient's decision or lack of a decision generates autonomy concerns if but only if the person has DMC.[11]

This is, on its face, a plausible way of operationalizing the Standard Story. It provides a more complete explanation of the intuitions in the cases of Jack and Jill—because Jill has DMC and Jack doesn't, respect for autonomy demands that Jill's decision has authority but makes no such demand of Jack's. And it's buttressed by the apparent overlap between DMC and the concept of autonomy—very plausibly self-governance requires at least some minimal levels of understanding, appreciation, and reasons-responsiveness.[12]

Despite this initial plausibility, though, the Standard Story and the Gatekeeping Model are worryingly ambiguous—or at least worryingly incomplete. Appeals to autonomy implicate at least two distinct ethical concerns. And varying which concern(s) DMC matters for yields different ways of understanding the Standard Story and the Gatekeeping Model.

4. Two Autonomy Concerns

To begin to get a grip on the distinction between the two autonomy concerns, consider the following pair of cases, adapted from Enoch (2017, pp. 6–7, 14):

> **Pat**. Pat arrives at your emergency room unconscious and in need of a blood transfusion to save her life. Blood transfusions typically require consent. And Pat isn't consenting—she's unconscious. Were she conscious, though, she would consent—she wants to keep living and has no qualms with invasive medical procedures.
>
> **Chris**. Chris arrives at your emergency room unconscious and in need of a blood transfusion to save his life. Blood transfusions typically require consent. And Chris

---

[11] See also Faden and Beauchamp (1986, Chapter 8) who explicitly label DMC a "gatekeeping concept."

A related, but nevertheless distinct, reading of the Standard Story takes DMC to determine when the reasons to respect autonomy *outweigh* reasons to protect or promote well-being (see Appelbaum, 2007; Appelbaum & Grisso, 1988; Buchanan & Brock, 1989). On this telling, the decisions (or their absence) of incapacitated patients might generate autonomy concerns, but such concerns are trumped by well-being concerns. This view is more nuanced, and hence more complicated, than the Gatekeeping Model I consider here. For the sake of simplicity, I'll set it aside and just note that many of the criticisms I level against the Gatekeeping Model in §5 apply, with some modification, to this alternative view.

[12] Beauchamp and Childress (2013, p. 116) make something like this point.

But note that my "minimal" hedge is important, here. As I'll argue in §5.2, on the most natural understanding of autonomy, the requirements of autonomy are *weaker* than those of DMC.

> isn't consenting—he's unconscious. Moreover, even if he were conscious, he still wouldn't consent—he's a Christian Scientist and his convictions forbid transfusions.

With Enoch, I'm inclined to think that we should give Pat the transfusion. I'm also inclined to think that we should *not* give Chris the transfusion.[13] And, in both cases, it seems that what's driving the verdict about what we should do is concern for the relevant person's autonomy.[14] Each is an autonomous agent and, as such, has the capacity to govern their life in accordance with their deep commitments—their firmly held beliefs, deeply held values, important aims, strong intentions, etc.[15] Since we're in a position to ensure that their lives in fact go in a way that accords with their deep commitments, respecting autonomy calls for giving Pat the transfusion and withholding it from Chris.

Now consider another pair of cases:[16]

> **Sally**. Sally is running late for a noon talk. It's looking like she won't arrive in time to pick up a provided lunch, so you decide to grab one for her. You know that she recently decided to go on a low-carb diet. The options are a salad or a hoagie. You text to ask what she wants. She responds: "the salad."
>
> **Hugo**. Hugo is running late for a noon talk. It's looking like he won't arrive in time to pick up a provided lunch. So you decide to grab one for him. You know that he recently decided to go on a low-carb diet. The options are a salad or a hoagie. You text to ask what he wants. He responds: "the hoagie."

---

[13] Of course, we'd need more details to vindicate this intuition. And, even then—depending on how strong you think our reasons to save Chris's life are—the intuition might not be very strong. These qualms can be set aside, though. To make the point, all I really need is for there to be a normative difference between the two cases. And certainly there is one: we at least have *more reason* to give Pat the transfusion than we do to give Chris the transfusion, even if we should ultimately give it to them both. So if you're uneasy about the judgment that we shouldn't give Chris the transfusion, focus instead on this difference in reasons, and the rest of the discussion may proceed unperturbed.

[14] You might disagree. You might think that it's concern for their *well-being* that's driving these verdicts. I don't think that's right. Consider someone whose deepest commitments are self-sacrificing. It seems to me that we have reason to ensure that their life goes in a way that's consistent with these commitments, but it doesn't seem to me that this reason is best explained by concern for their well-being. There's much more to say on this score, though, and a full investigation would take us too far afield. So let me set it aside by noting that whether (what I'll soon call) authenticity is best characterized as an autonomy concern or a well-being concern—though quite relevant to how this project is framed—it's not especially relevant to my general account of why DMC matters or to my criticisms of alternative accounts. So if you're tempted by the thought that our concern for Chris is a concern for well-being, then read future mentions of "authenticity" as appeals to an important well-being concern implicated in our use of DMC rather than an important autonomy concern.

[15] In what follows I'll use 'deep commitments' as a catch-all for these aspects of an autonomous agent.

[16] While these cases aren't straightforwardly adapted from Enoch, they're inspired by his discussion of his daughter passing him the salt (Enoch, 2017, pp. 31–32).

Here, I'm inclined to think that we should grab Sally the salad. I'm also inclined to think that we should get Hugo the hoagie.[17] And, in both cases, it seems that what's driving the verdict about what we should do is concern for the relevant person's autonomy. But importantly, this is a quite different concern from that which drove the verdicts in the cases of Pat and Chris. There, we were just concerned with their deep commitments. Here, however, we can stipulate that there's *no difference* between the deep commitments of Sally and Hugo, and, still, autonomy seems to demand disparate treatment. Our concern for autonomy, then, isn't just a concern that a person's life goes in a way that's consistent with her deep commitments; it's also a concern that she exercises control over that which is properly hers to control. Since we're in a position to ensure that Sally and Hugo in fact exercise control over that which is properly theirs to control, respecting autonomy calls for grabbing Sally the salad and Hugo the hoagie.

What this shows is that there are two morally important facets of an autonomous agent: an autonomous agent is capable of *governing* her life in *accordance with her deep commitments*. And, given this, there are two ways in which we can respect an agent's autonomy: we can ensure that her life in fact goes in a way that accords with her deep commitments, and we can ensure that she's the one that does the governing. This distinction has been marked by different philosophers in different ways. Here, I'll adopt the labels 'authenticity' and 'sovereignty'.[18]

---

[17] Here again you may not share the intuition about what we should do for Hugo. But, also again, this concern can be set aside. To make the point, all that's necessary is that we have more reason to get Sally the salad than we do Hugo. And this seems clear even if, all things considered, we shouldn't get Hugo the hoagie.

[18] I adopt the former from Brudney and Lantos (2011) and the latter from Enoch (2017). What Enoch calls 'sovereignty', Brudney and Lantos call 'agency', and what Brudney and Lantos call 'authenticity', Enoch calls 'nonalienation.'

Authenticity: a choice or outcome is authentic for a person to the extent that it reflects or accords with her deep commitments (her firm beliefs, strong desires, sacred values, stringent intentions, important aims, etc.).

Sovereignty: a person has sovereignty over a choice or outcome to the extent that it is properly hers to control (i.e. is within her *sovereign domain*).[19]

So, concern for a person's autonomy sometimes manifests in concern for authenticity—concern that her life goes in a way that's consistent with her deep commitments—and sometimes manifests in concern for sovereignty—concern that she exercises control over that which is properly hers to control. And this makes the Standard Story start to look murky. Sure, DMC matters because autonomy matters. But does it matter because authenticity matters? Because sovereignty matters? Or because both matter?

5. Why DMC Matters: A More Complicated Story

5.1 DMC matters because authenticity matters

Consider first the possibility that DMC matters because authenticity matters. This looks plausible on its face. Authenticity, one might think, is the more *fundamental* autonomy concern, and sovereignty demands respect only insofar as respecting it allows autonomous agents to better pursue their deep commitments.[20] And this suggests that, on the Standard Story, the moral importance of DMC must ultimately be grounded in concern for authenticity.

But reflection reveals this can't be right. Clearly, incapacitated patients can (and often do) have deep commitments. Chris, from above, is a case in point. He clearly lacks DMC when he arrives at the emergency room—he's unconscious. But of course this fact about Chris does not preclude him

[19] Importantly—and as will become clear in what follows—to claim that a person has sovereignty over some decision is *not* to claim that she ought to control it *all things considered.* That a person has sovereignty over some decision is a very important consideration, but it's just one among many. And it's possible that in some cases it will be outweighed, such that a person has sovereignty over a decision—i.e. considering her autonomy alone, it is properly hers to control—but, all things considered, she should *not* control it.

[20] This is one way of reading Brudney and Lantos (2011).

from being deeply dedicated to Christian Science or from having particular preferences regarding blood transfusions. Clearly, authenticity concerns arise for Chris despite his being incapacitated.

In response, an advocate of this version of the Standard Story might counter that while Chris lacks DMC at the time he arrives in the emergency room, there's a more general sense in which he plausibly has (or at least had) DMC. If Chris's incapacity is recent or temporary, then one might maintain that it's in virtue of Chris's previous or future capacity that authenticity concerns arise. But while such a response can countenance authenticity concerns in cases of recent or temporary incapacity, it can't accommodate authenticity concerns in cases in which a patient chronically lacks DMC. But in at least some such cases, authenticity concerns clearly arise. If Chris came into the emergency room conscious but with a learning disability that rendered him incapacitated, his deep commitment to Christian Science and strong preference against transfusion would raise authenticity concerns and be quite relevant to how we should treat him despite his lack of DMC. And Chris (so imagined) isn't alone. With little effort, we could come up with countless other cases in which a person chronically lacks DMC but nevertheless has deep commitments. This is because the abilities constitutive of DMC just aren't necessary, or even tightly correlated with, the capacity for deep commitments. And if we hold fixed a patient's deep commitments, then whether that patient has DMC makes no difference to whether any particular option conforms to or conflicts with those commitments. So to the extent that we're concerned just with autonomy-as-authenticity, DMC is beside the point.

Finally, an advocate might rejoin that while a lack of DMC doesn't entail, or even strongly indicate, the *absence* of deep commitments, it does render them inaccessible and hence irrelevant. This is something like the thought underlying the 2007 DC Circuit Court ruling in *Doe Tarlow v. District of Columbia*, that, for patients who have never had DMC, their "true wishes… 'are unknown and cannot be ascertained.'" But while it might be true that a lack of DMC can be correlated with an inability to explicitly express one's deep commitments, this isn't enough to render authenticity considerations irrelevant. First, there are a number of *other* sources from which we might glean a person's deep

commitments other than her explicit expression of them.[21] Second, this correlation is far from robust. The abilities relevant to having and expressing commitments might overlap to an extent with those constitutive of DMC. But a person who lacks one or more such abilities might nevertheless have no trouble clearly and coherently telling us what she cares about.[22] So even if a lack of DMC *sometimes* forecloses *one* way of knowing about a patient's deep commitments, this clearly isn't enough to vindicate this version of the Standard Story.

5.2 DMC matters because sovereignty matters

Next, consider the possibility that DMC matters because sovereignty matters. This too looks initially plausible. After all, to have sovereignty over some choice or outcome *just is* for it to be properly yours to control. So sovereignty might look well positioned to explain why DMC matters to whether, all things considered, a patient's decision should determine her treatment.

Ultimately, though, this proposal is beset by concerns similar to those that plague the previous one. Recall that a person might make a decision despite *lacking* DMC. Given this, it's puzzling as to why whether a choice or outcome is within a patient's sovereign domain should depend on her having or lacking DMC.

Consider, for example, the following patient.

> **Dilara**. Dilara is a patient with moderate dementia. Moderate, but enough to render her incapacitated regarding medical decisions—she's incapable of understanding even the most basic details of a medical procedure, or of appreciating the long-term effects of her ailments, or of reasoning well about the available medical options. She develops a kidney condition that requires dialysis. But she adamantly refuses: "I'd like my blood to stay in my body," she insists. And there's no convincing her otherwise.

---

[21] In fact, on many accounts, special access to these sources of information about a person's deep commitments is what justifies the use of surrogates to make decisions on behalf of incapacitated patients (see Kelly et al., 2012).

[22] For more on these points, see Howard and Wendler (2020). It's their paper that alerted me to the *Tarlow* case.

Here, Dilara has made a genuine decision to forego dialysis.[23] And, assuming (as is plausible) that whether your blood stays in your body is an outcome that is properly yours to control, it seems that to give Dilara dialysis would be to *fail* to respect her sovereignty, and to *not* give her dialysis would be to respect it. And this seems so whether or not Dilara has DMC.

More generally, if what grounds our reason to do as a patient decides is concern for sovereignty, then it's not clear why these grounds would reliably evaporate in the case of an incapacitated patient. If we hold fixed a patient's decision, then whether that patient has DMC makes no difference to whether any particular option is consistent or inconsistent with that decision. So to the extent that we're concerned just with autonomy-as-sovereignty, DMC is again beside the point.

Perhaps, though, this isn't the most charitable reading of this version of the Standard Story. In response, an advocate might insist that I've missed the point: the reason that sovereignty considerations don't arise for patients who lack DMC is that such patients' decisions *aren't autonomous* and hence don't generate autonomy concerns. This is to attempt to vindicate this version of the Standard Story by *equating* autonomy and DMC and thereby insisting that it's *not possible* for someone to *autonomously* decide without DMC.[24] The thought is that a person exercises control over a choice or outcome only if they make an autonomous decision about that choice or outcome; and since being autonomous—i.e. having DMC—is a prerequisite for making autonomous decisions, when a patient lacks DMC, respecting her sovereignty *isn't an option* because nothing she does qualifies as her exercising control. So, on this view, if a person lacks DMC, she is incapable of making an autonomous decision, and hence her (non-autonomous) decisions don't raise sovereignty concerns because they don't qualify as exercises of her autonomy.

---

[23] Recall that to make a decision doesn't take much—it's just to form an intention or to settle a practical question or to make up one's mind about what to do (see §2). And so I don't think there should be any trouble with supposing that Dilara's refusal constitutes a genuine decision. But perhaps you think that Dilara's decision isn't of the right *sort* because it's not an *autonomous* decision. I'll consider that line just below.

[24] For two examples of this approach, see Burch (2017, p. 3) and Faden and Beauchamp (1986, p. 288). It's also at least one way of reading the Belmont Report (1979).

This is, to my eyes, a forceful response. But I don't think it can be made to work. In §2, I defended the possibility of decisions without DMC. The reading of the Standard Story currently on offer maintains that all such decisions are *non-autonomous*. And that just isn't so. Recall that to be autonomous is to be self-governing and to act autonomously is to exercise one's capacity for self-governance. To *fail* to act autonomously, then, is to *fail to exercise this capacity*—for what you do to float freely of your capacity to determine how you will act.[25] Of course, the details about what exactly makes for an autonomous decision are a matter of philosophical dispute. But we needn't get bogged down in them to appreciate this point: the requirements of autonomous action are often *weaker* than those of DMC. One might lack the ability to understand or appreciate or reason or communicate—to the extents required by DMC—while still exercising one's general capacity for self-governance; while manifesting the kind of control over one's actions in light of one's commitments that's constitutive of autonomous action.[26] A delusional patient is a paradigm case of lacking the sort of understanding or appreciation required for DMC. But delusions—even systemic delusions—aren't typically enough to undermine a person's capacity to determine how she'll act in the autonomy-constituting way.[27] Patients with dementia are often cited as illustrative examples of failing to meet DMC's reasoning requirement. But only in the most extreme cases is it plausible that a demented patient's reasoning capacities are so impaired as to render her incapable of self-governance.

To illustrate, consider again Dilara. Imagine her insisting. Imagine pressing her on her decision—on the false beliefs and flawed reasoning that underlie it—and imagine her continuing to insist. Now,

---

[25] My characterization of failure to act autonomously is reminiscent of Buss and Westlund (2018, sec. 1, emphasis added): "If an agent fails to govern herself when she acts, this must be because *what she does is independent of her power to determine how she will act*." See their piece for an especially helpful overview of accounts of autonomous action and to get a sense of the ways in which the requirements of autonomous action are thin as compared to those of DMC.

[26] While, as mentioned at the end of §3, it's plausible that *some* level of understanding of (and appreciation of and reasoning about) what one is doing is necessary for an action to qualify as autonomous, it's typically not the levels required by DMC. It's also worth pausing to note and explain my 'often' and 'typically' hedges. Since the requirements of DMC fluctuate with the stakes of the relevant decision (see footnote 3), it's possible that for some very low-stakes decisions, the requirements of autonomy are *not* weaker than those of DMC. (Thanks to an anonymous reviewer for pointing this out.)

[27] In fact, part of the reason we reasonably *worry* about a delusional patient making decisions for herself is that we expect her to be governing herself in accordance with her delusions.

I admit that the fact that she lacks DMC is relevant to how we should ultimately treat her (and I'll attempt to say exactly how it's relevant in the next section). But does it render her decision *not autonomous*? I don't see how it could. Clearly, Dilara has the *general* capacity to self-govern—she makes decisions in light of her commitments about what to eat and what to wear and how to spend her time, just as you and I do. And, just as clearly—or so it seems to me—her decision to refuse dialysis is an exercise of this capacity. It's not the expression of an addiction or an uncontrolled outburst. It's an instance of Dilara's general capacity to self-govern—an instance of her commitments combining in the appropriate way to determine what she does. Admittedly, I made up Dilara. But there's nothing fanciful about her case. And countless real-world examples could illustrate the point: intuitively autonomous decisions are possible for patients that lack DMC.[28]

This possibility of autonomous decision-making without DMC becomes especially striking when we note two previously mentioned features of DMC as its assessed and measured. First (as mentioned in footnote 3), it's typically thought that the ability thresholds required for DMC go up as the stakes of the choice go up.[29] The thought is that when the stakes are high, DMC requires proportionately greater abilities to understand, appreciate, reason, and communicate than are required when the stakes are low. But no one thinks that whether someone is able to make an autonomous decision fluctuates with what's at stake. Second, it's just not particularly plausible that the full 2.8% of healthy elderly

---

[28] Importantly, this isn't to suggest that every person is autonomous, nor is it to suggest that every decision of an autonomous person is autonomous. Again, it's very plausible that self-governance requires at least some minimal levels of understanding, appreciation, or reasons-responsiveness (in addition, perhaps, to some other abilities). So, for instance, if Dilara were wholly incapable of means-end reasoning or unable to form intentions or plans, she might fail to qualify as autonomous. Likewise, an autonomous person's *decision* will not itself qualify as autonomous if that decision fails to be a manifestation of the person's general capacity for self-governance. So, for instance, if Dilara's decision were driven by a compulsion and unconnected to her broader beliefs and commitments, it might be non-autonomous. Of course, a complete accounting of exactly when agents and their decisions are and are not autonomous is well beyond the scope of the current project. And, here, I'd like to remain (for the most part) neutral among the competing accounts. But it's nonetheless important to note that while the requirements of autonomous decision making are (often) weaker than those of DMC, they are not trivially met. (My thanks to an anonymous reviewer for pressing me to clarify this point.)

[29] For influential discussions, see Appelbaum (2007) and Buchanan and Brock (1989, Chapter 1). Notably, some even argue that the thresholds fluctuate with the *content of a patient's decision*—such that when choosing between risky option A and less risky option B, the level of understanding, etc. required for DMC is higher if she chooses A than if she chooses B. See, for example, Wilks (1999).

adults who lack DMC (or 20% of patients with mild cognitive impairment, or 68% of learning disabled patients, etc.) are also incapable of making autonomous decisions in the relevant domain, whatever it is.[30] It's much more plausible to think that at least some of these individuals retain the ability to autonomously decide despite lacking DMC.

One last consideration about this proposal is worth addressing: it's *revisionary* in one way or another, and either way is problematic. Since, on the most natural understanding of autonomy, its requirements are (often) *weaker* than those of DMC, any attempt to equate autonomy and DMC must either beef up the requirements of autonomy or tone down those of DMC.[31] The former route insists that autonomy is *more* than the mere capacity to govern oneself according to one's commitments and that it involves understanding, appreciation, reasoning, and communication to the degrees required for DMC (again, as it's assessed and measured). But if we engineer our concept of autonomy in this way, then it will be ill-suited to explain the important role of *respect for autonomy* in normative theorizing and applied ethics. It's widely agreed that we have at least a strong pro tanto reason to respect the autonomy of others. And it's also widely agreed that this reason is ubiquitous—nearly everyone has autonomy of the sort that demands respect, and the reason to respect a person's autonomy persists even when her action is uninformed (because she can't understand), imprudent (because she can't appreciate), irrational (because she can't reason well), or otherwise mistaken. But if autonomy is thought to require the abilities constitutive of DMC, then our reason to respect autonomy is much more circumscribed—we lack such a reason whenever a person can't understand or appreciate some of the relevant information or reason well about some of her options; in fact, we lack such a reason for 2.8% of healthy elderly adults. Alternatively, the latter route—opting to equate autonomy and

[30] As mentioned in §2, referencing Sessums (2011).

[31] Admittedly, there are more robust accounts of autonomy on offer in the literature—see, for example, Dworkin (1993). But, for precisely the reasons canvassed here in this section, these accounts strike me as ill-suited to explain autonomy of the sort regularly exhibited by actual agents. They are instead—to my eyes anyway—most plausible when interpreted as accounts of an *ideal* of autonomy, which is only indirectly relevant to the project at hand. See Benn (1976) for a helpful discussion of autonomy-as-an-ideal.

DMC by weakening the requirements of DMC—would make it uncontroversial that DMC-*so-revised* matters because autonomy matters, but would thereby abandon the project of defending the Standard Story, which maintains that DMC-*as-conventionally-understood* matters because autonomy matters. And insisting that some alternatively conceived DMC matters because autonomy matters is just to engage in a different project. I'll say a bit more about this option in §5. For now, the upshot is that the version of the Standard Story that equates DMC and autonomy ultimately can't get off the ground.

Finally, one might try to secure this reading of the Standard Story—that DMC matters because sovereignty matters—by maintaining that while it's true that respecting a patient's sovereignty requires ensuring that she exercises control over that which is properly hers to control, *whether* a medical outcome is properly a patient's to control is determined by whether she has DMC. On this interpretation, then, DMC fixes the boundaries of one's sovereign domain such that, in medical contexts, when a patient lacks DMC, sovereignty considerations don't arise because choices in such contexts are not properly hers to control.

This proposal is problematic in at least three respects. First, it's simply counterintuitive to insist that sovereignty considerations just don't arise for a patient like Dilara or that whether her blood stays in her body is not hers to control. I'm certainly open to the possibility that, all things considered, we ought to give Dilara dialysis despite her protestations (for the sake of her well-being, or for her deep commitments, etc.). But surely we have a *reason* to *not* give her the dialysis that we wouldn't have were she to happily agree to it. And this reason seems to be grounded in Dilara's autonomy. This suggest that sovereignty matters in Dilara's case, even if it's ultimately outweighed and despite her lack of DMC. Second, the abilities constitutive of DMC seem to be of the wrong sort to determine whether something is within a person's sovereign domain. In an influential discussion of sovereignty, Feinberg (1989, p. 55) maintains that "My personal [sovereign] domain… consists of my body, privacy, landed and chattel property, and at least the vital life-decisions, perhaps among other things." And there's no reason to think he was just talking about himself. What's notable about his list is that

each item characterizes a sphere of influence that makes no reference to features of the individual who may or may not be exercising control. The very plausible thought underlying this is that what determines whether a domain is properly yours to control are features *of that domain*—of what it's about or of its type—and *not* features *of you*. In other words, once we determine the relevant domain, we can know whether your decision within it has weight—i.e. whether it's properly yours to control—without knowing whether or not you have DMC. Third, in many non-medical contexts, we don't maintain that sovereignty matters only for persons with the abilities constitutive of DMC. When making investments, pursuing relationships, having children, choosing careers, picking hobbies, finding religion, etc. we generally do not think that a person loses sovereignty over these choices if she is incapable of understanding some of the relevant information, or appreciating some of the relevant circumstances, or reasoning about some of the relevant options. Given this, to vindicate this version of the Standard Story, we'd need some set of considerations that are salient in the medical context that render sovereignty considerations irrelevant or less weighty in incapacitated patients. And I don't see what those considerations could be. In fact, given that medical decisions typically concern one's own body, mind, and general well-being, it seems that, if anything, the medical context would be one in which sovereignty considerations are especially weighty and difficult to discharge.

5.3 DMC matters because harmony matters

So, it's not particularly plausible that DMC matters just because authenticity matters or just because sovereignty matters. Perhaps instead, then, it matters because *both* authenticity and sovereignty matter. Here, there's a compelling story to tell. The abilities that constitute DMC are abilities that facilitate controlling one's life in such a way that it *actually accords* to one's deep commitments. A person who fully understands the relevant information, wholly appreciates how it applies to her, carefully reasons about the available options, and clearly communicates her choice is well—and perhaps *best*—positioned to choose whatever option is *in fact* most consistent with her deep commitments.

For this reason, a capacitated person is more likely to agree to that which is *actually consistent* with her deep commitments and reject that which is *actually inconsistent* with her deep commitments than a similarly situated person who lacks DMC. In this sense DMC indicates a *harmony* between respecting authenticity and respecting sovereignty. Presented with a patient whose choice of treatment has been made clear but whose deep commitments are unknown, if that patient has DMC, then, absent countervailing evidence, we have good reason to think that her choice is one that conforms to her deep commitments. Similarly, presented with a patient whose deep commitments are known but whose choice is unclear—suppose, for instance, we have good evidence *that* she made a decision, but no way of discerning the *content* of her decision—if that patient has DMC, then, absent countervailing evidence, we have good reason to think that the option that best conforms to her deep commitments is the one she chose. In a patient without DMC, we're not in a position to make these inferences from choice to deep commitments and deep commitments to choice. And, to the extent that we care about both sovereignty and authenticity, this is why an incapacitated patient's decision should not necessarily determine her treatment—we must first ensure that her choice is consonant with her deep commitments, and if it's not, we must weigh the value of respecting her sovereignty against the value of respecting her authenticity.

To make all this a bit more concrete, consider again Jack and Jill (from §1). Recall that Jill has DMC and Jack does not and that both opted for the chemo over the palliative care. On the account of why DMC matters on offer here, the fact that Jill has DMC indicates that her choice is consistent with her deep commitments. So, we have good reason to believe that respecting Jill's sovereignty—i.e. doing as she decides—is also to respect her authenticity—i.e. to do what's most consistent with her deep commitments. In contrast, since Jack lacks DMC, we lack this source of evidence about which treatment option best accords with what he cares about. And hence, even when considering his autonomy alone, it's not clear what we should do with him—giving him the chemo would grant him control over that which is properly his to control, but it may run afoul of his deep commitments. All

this provides a yet deeper explanation of the intuitions about Jack's and Jill's respective cases. Jill's decision should determine her treatment because she has DMC, and this fact matters because autonomy matters—specifically, because it indicates harmony between the two autonomy concerns. Jack's decision should not necessarily determine his treatment because, for him, to respect his autonomy-as-sovereignty might be to violate his autonomy-as-authenticity.

This vindicates the Standard Story about why DMC matters. But it introduces a twist. DMC does matter to when a patient's decision should determine her treatment, and it does matter in this way because autonomy matters. But the role DMC is playing is an *evidential* one, not a *moral* one. DMC is not *itself* a part of the moral grounds for why we should do as a capacitated patient decides (in those cases in which we should) or of the moral grounds for why we should not do as an incapacitated patient decides (in those cases in which we shouldn't). Instead, the moral grounds are exhausted by facts about authenticity and sovereignty alongside other morally relevant considerations (like the patient's well-being, the family's wishes, etc.), and DMC matters in that it serves as *evidence* that the moral grounds are a certain way—specifically as evidence that the reasons generated by sovereignty and the reasons generated by authenticity point in the same direction.

6. A Beacon, not a Gate

So, the most plausible account of why DMC matters in terms of autonomy is that the abilities constitutive of DMC make a person's decisions *reliable indicators* of her deep commitments (and vice versa). Again, this account vindicates (with a twist) the Standard Story about why DMC matters. But, crucially, it doesn't vindicate the Gatekeeping Model. While a capacitated patient's decisions are *more likely* to be consistent with her deep commitments than an incapacitated patient's decisions are to be with hers, *this indication is defeasible*. It's possible for a capacitated patient to choose *against* her deep commitments just as it's possible for an incapacitated patient to choose *in accordance* with hers. And this makes trouble for the Gatekeeping Model.

To see, consider one last time Jack and Jill. This time, though, further suppose that we're *certain* of both Jack's and Jill's respective commitments and we're *certain* that the chemo option conflicts with each. These are, admittedly, rather odd things to suppose. First, we're not often (and, plausibly, not ever) certain of the commitments of others. Second, we're supposing that, despite having DMC, Jill has made a *mistake*; she chose poorly. In less fanciful circumstances, the fact that Jill has opted for the chemo would make us doubt whether we really know her commitments or whether she really has DMC. But here—in order to make more visible the Gatekeeping Model's problems—I'm *stipulating* that 1) Jill has DMC, 2) she has chosen the chemo, and 3) the chemo conflicts with her deep commitments. And while this is an odd combination of facts, it's not an impossible one. Jill's decision is difficult, and we know that even informed, insightful, and otherwise rational individuals are subject to weakness of will, implicit biases, and other mistakes of reasoning. So it's certainly possible that Jill chose against her deep commitments despite being able to understand the relevant information, appreciate that it applies to her, reason about her options, and communicate her choice. And that's just what we're supposing: that she, along with Jack, has chosen the chemo, and has done so *despite* the fact that it conflicts with her commitments.

Again, Jill has DMC and Jack does not. And, given this, the Gatekeeping Model insists that only Jill's decision generates autonomy concerns. But from the perspective of autonomy, Jack and Jill are identical—concern for sovereignty demands that we do as each patient chooses and give each the chemo; respect for authenticity demands that we do what's consistent with each patient's deep commitments and give each palliative care.[32] Here, these two demands conflict. And however the tradeoff should be made, as far as autonomy is concerned, it should be made identically for each. Given that we're certain about the relevant autonomy considerations—the content of Jack's and Jill's

[32] Of course, this assumes that Jack's impairments, though enough to render him incapacitated, aren't so severe as to render him non-autonomous. In §5.2, I argued for this possibility. Here, I just stipulate that Jack is an instance of it.

respective choices and the character of their respective commitments—the fact that one has DMC and the other does not is beside the point.

If we hold onto the Standard Story, then, the Gatekeeping Model fails on two scores. First, it mistakenly insists that autonomy concerns arise only from the decisions of capacitated patients and thereby fails to grapple with important tradeoffs between the two autonomy concerns—both in incapacitated patients whose sovereignty is ignored for the sake of their deep commitments and in capacitated patients whose deep commitments are ignored for the sake of their sovereignty. Second, it suggests differential treatment in cases—like Jack's and Jill's—in which there's no morally relevant difference to justify it.

More generally, while I'm not sure exactly how tradeoffs between authenticity and sovereignty should be made, *if* we think that a capacitated patient's decision should, in the normal range of cases, have authority *even if* we know that her decision conflicts with her deep commitments, then I don't see how, on the basis of concern for autonomy, we can justify not giving authority to an incapacitated (yet autonomous) patient's decision in similar circumstances. And, likewise, if we think that an incapacitated (yet autonomous) patient's decision should *not* be given authority when we know that it's sufficiently contrary to her deep commitments, then I don't see how, on the basis of concern for autonomy, we can justify giving authority to a capacitated patient's decision in similar circumstances. In any case, if we insist on telling the Standard Story about why DMC matters, we should stop treating it as a gatekeeper for autonomy.

The lesson of all this is that the Standard Story about why DMC matters and the Gatekeeping Model form an unhappy marriage. Given a plausible account of the relevant autonomy concerns, the Gatekeeping Model is an inappropriate way of determining when autonomy considerations are and aren't relevant. This leaves us with a few ways out. We could keep (something like) the Gatekeeping Model and reject the Standard Story—justifying our use of DMC by appealing to something other than (or in addition to) autonomy. Or we could keep the Gatekeeping Model and the Standard Story

and reconceptualize DMC so that it better tracks autonomy. Or we could preserve our notion of DMC and the Standard Story about why it matters and reject the Gatekeeping Model. My sympathies lie with the last option. But I think there's something to be said for each. So I'll close with some mostly suggestive remarks about where we might go from here.

I've argued that the most plausible reading of the Standard Story about why DMC matters is inconsistent with the Gatekeeping Model. In light of this, one way to preserve (something like) the Gatekeeping Model is to tell a different story about why DMC matters. I'm sure there are many candidates. But one strikes me as especially appealing. How to treat patients in the course of medical research and clinical care is a real-life practical problem. And we need some (more or less) *objective, applicable* criteria for determining when a patient's decision should determine her treatment. DMC gives us that—a well-established, validated measure of abilities that are relevant (evidentially and defeasibly, I've argued) to whether a patient's decision should determine her treatment. A plausible reason to *ignore* relevant autonomy considerations when a person lacks DMC, then, is that to attempt to be responsive to such considerations in incapacitated patients would be, practically speaking, unworkable. The thought is that the value of having a workable policy is worth taking on the disvalue of ignoring important autonomy considerations for a nonnegligible subset of incapacitated patients.[33] I'm skeptical. It seems to me that we could construct workable policies that capture the nuances of the two autonomy concerns. But I've no argument to support this sense, and I'm far from confident. So I take it to be a live possibility that, as a matter of policy, perhaps we *should treat* DMC as a gatekeeper for autonomy even though it doesn't perfectly capture when and why autonomy concerns are relevant. But even if this is so, such a policy should be built upon solid moral grounds, and we should be explicit about its costs and about the justification for why those costs are worth bearing. The first step towards

[33] This strategy is suggested by Buchanan and Brock (1989) when they justify the use of (something like) the Gatekeeping Model on the grounds that it will minimize occurrences of certain types of errors (specifically, errors of violating autonomy and of failing to protect the vulnerable).

doing so is to get clear about the ways in which autonomy considerations are and aren't relevant when a patient has or lacks the abilities constitutive of DMC, irrespective of the broader policy considerations. So, even if this story is ultimately right—if DMC matters because it's an element of the best *policy* for determining when a patient's decision should determine her treatment—then one way of reading the preceding project is as an attempt to take this first step.

The other way to preserve the Gatekeeping Model is to keep telling the Standard Story about why DMC matters but *reconceptualize* DMC so that it better tracks autonomy. I mentioned this suggestion in §5.2. It strikes me as a plausible proposal insofar as it's incontrovertible that autonomy considerations arise regarding a patient's decision if and only if that patient's decision is autonomous. So equating DMC and autonomy would clearly vindicate both the Standard Story and the Gatekeeping Model. And if DMC was *intended* to track autonomy to begin with, then perhaps this is the tack we should take.[34] But it's worth noting the costs of this strategy. As I've argued, DMC *as currently conceived* matters to whether a patient's decision should determine her treatment. And as I mentioned just above, we have well-established and validated measures for the abilities that constitute DMC. I know of no account of autonomy that comes with a well-established and validated measure. So reconceptualizing DMC in the manner suggested risks abandoning a useful tool in determining how we should treat patients without any clear prospect of replacing it with a similarly well-functioning one.

Finally, we might opt to reject the Gatekeeping Model. This option strikes me as the most attractive because, again, DMC matters as currently conceived and, as the Standard Story goes, it matters because autonomy matters. As I've argued, DMC indicates a harmony between respect for authenticity and respect for sovereignty. And this is not a trivial role for it to play. The demands of autonomy are often difficult to discern, and DMC can be a beacon in the fog. It's something for which

[34] For what it's worth, Hawkins and Charland (2020, sec. 2) note that DMC was *not* conceptualized to track any philosophical understanding of autonomy.

we have well-developed measures, and it indicates something important—namely, that to do as a patient decides is also to do what she really wants. But capturing this way in which DMC matters requires abandoning the Gatekeeping Model. This amounts to both a defense of the importance of DMC as currently conceived and a call for substantial revision to how it's used in medical research and clinical care. If the preceding project succeeds, then we can't just ignore autonomy considerations when treating incapacitated patients, and we can't just assume that respecting the autonomy of capacitated patients invariably demands that we do as they decide. Given the distinction between authenticity and sovereignty, matters are more complicated than that. Whether a patient's decision should have authority turns out to depend not just on whether she's made an autonomous decision about something that's properly hers to control, but also on the extent to which that option conforms to her deep commitments. Determining whether a patient has DMC can *help* us discern autonomy's demands, but it can't settle the question all on its own. Treating DMC as a gatekeeper for autonomy fails to grapple with the important tradeoffs that, in some cases, must be made between authenticity and sovereignty. Recognizing DMC's appropriate role, then, is to recognize the extent to which respecting autonomy is more difficult and more complicated than bioethicists often suppose.

References

Appelbaum, P. S. (2007). Assessment of Patients' Competence to Consent to Treatment. *New England Journal of Medicine*, *357*(18), 1834–1840. https://doi.org/10.1056/NEJMcp074045

Appelbaum, P. S., & Grisso, T. (1988). Assessing Patients' Capacities to Consent to Treatment. *New England Journal of Medicine*, *319*(25), 1635–1638. https://doi.org/10.1056/NEJM198812223192504

Appelbaum, P. S., & Grisso, T. (1995). The MacArthur Treatment Competence Study I: Mental illness and competence to consent to treatment. *Law and Human Behavior*, *19*(2), 105–126. https://doi.org/10.1007/BF01499321

Beauchamp, T. L., & Childress, J. F. (2013). *Principles of Biomedical Ethics* (7th ed.). Oxford University Press.

Benn, S. I. (1976). VI—Freedom, Autonomy and the Concept of a Person. *Proceedings of the Aristotelian Society*, *76*(1), 109–130. https://doi.org/10.1093/aristotelian/76.1.109

Boyle, M. (2011). 'Making up Your Mind' and the Activity of Reason. *Philosophers' Imprint*, *11*(17), 1–24.

Brudney, D., & Lantos, J. (2011). Agency and authenticity: Which value grounds patient choice? *Theoretical Medicine and Bioethics*, *32*(4), 217–227. https://doi.org/10.1007/s11017-011-9180-2

Buchanan, A., & Brock, D. W. (1989). *Deciding for Others: The Ethics of Surrogate Decision Making*. Cambridge University Press.

Burch, M. (2017). Autonomy, Respect, and the Rights of Persons with Disabilities in Crisis. *Journal of Applied Philosophy*, *34*(3), 389–402.

Buss, S., & Westlund, A. (2018). Personal Autonomy. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Spring 201). https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/

Dworkin, R. (1993). *Life's Dominion*. Harper Collins.

Enoch, D. (2017). Hypothetical Consent and the Value(s) of Autonomy. *Ethics*, *128*(1), 6–36. https://doi.org/10.1086/692939

Faden, R. R., & Beauchamp, T. L. (1986). *A History and Theory of Informed Consent*. Oxford University Press.

Feinberg, J. (1989). The Moral Limits of the Criminal Law Volume 3: Harm to Self. In *The Moral Limits of the Criminal Law Volume 3: Harm to Self*. Oxford University Press. https://doi.org/10.1093/0195059239.001.0001

Grisso, T., & Appelbaum, P. S. (1995). The MacArthur Treatment Competence Study. III: Abilities of patients to consent to psychiatric and medical treatments. *Law and Human Behavior*, *19*(2), 149–174. https://doi.org/10.1007/BF01499323

Grisso, T., Appelbaum, P. S., Mulvey, E. P., & Fletcher, K. (1995). The MacArthur Treatment Competence Study II: Measures of abilities related to competence to consent to treatment. *Law and Human Behavior*, *19*(2), 127–148. https://doi.org/10.1007/BF01499322

Hatherley, J. J. (2019). Is the exclusion of psychiatric patients from access to physician-assisted suicide discriminatory? *Journal of Medical Ethics*, *45*(12), 817–820. https://doi.org/10.1136/medethics-2019-105546

Hawkins, J., & Charland, L. C. (2020). Decision-Making Capacity. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2020). https://plato.stanford.edu/archives/fall2020/entries/decision-capacity

Hieronymi, P. (2009). Two kinds of agency. In L. O'Brien & M. Soteriou (Eds.), *Mental Actions* (pp. 138–162). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199225989.001.0001

Hindmarch, T., Hotopf, M., & Owen, G. S. (2013). Depression and decision-making capacity for treatment or research: a systematic review. *BMC Medical Ethics*, *14*(1), 54. https://doi.org/10.1186/1472-6939-14-54

Howard, D., & Wendler, D. (2020). Beyond Instrumental Value: Respecting The Will of Others and Deciding On Their Behalf. In A. Cureton & D. Wasserman (Eds.), *The Oxford Handbook of Philosophy and Disability*. Oxford University Press.

Kelly, B., Rid, A., & Wendler, D. (2012). Systematic Review: Individuals' Goals for Surrogate Decision-Making. *Journal of the American Geriatrics Society*, *60*(5), 884–895. https://doi.org/10.1111/j.1532-5415.2012.03937.x

Kim, S. Y. H. (2009). *Evaluation of Capacity to Consent to Treatment and Research*. Oxford University Press.

Mele, A. (2017). *Aspects of Agency: Decisions, Abilities, Explanations, and Free Will*. Oxford University Press.

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. (1979). *The Belmont Report*. https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report

Sessums, L. L., Zembrzuska, H., & Jackson, J. L. (2011). Does This Patient Have Medical Decision-Making Capacity? *JAMA*, *306*(4), 420. https://doi.org/10.1001/jama.2011.1023

Vellinga, A., Smit, J. H., van Leeuwen, E., van Tilburg, W., & Jonker, C. (2004). Instruments to assess decision-making capacity: an overview. *International Psychogeriatrics*, *16*(4), 397–419. https://doi.org/10.1017/S1041610204000808

Wilks, I. (1999). Asymmetrical competence. *Bioethics*, *13*(2), 154–159.