

RESEARCH ARTICLE

Health Policy

David Schmidtz

Presidential Chair of Moral Science, West Virginia University
Email: dschmidtz01@gmail.com

Abstract

In clinical trials involving experimental subjects who are also patients, what is supposed to become of the imperative to focus on the patient's best interest? A second set of policy questions concerns patients who want to die. Are there limits to the imperative to let patients choose for themselves? Is commodification a threat to autonomy? When, if ever, do costs and benefits become decisively important? Can we know what to count as a cost-effective preparation for the next pandemic? When we put procedures in place to protect against abuse, is there any way to prevent such measures from becoming bureaucratic obstacles to accomplishing anything at all?

Keywords: trolley problem; human subjects; cost-benefit analysis; consent; eugenics; killing versus letting die

Is basic research tragic?

The subtext of Lynn Jansen's lead essay for this issue is that the human condition is a challenge that we navigate with limited, sometimes tragically limited, information. Responding to our limits by gathering information and pushing our limits is characteristically human and often inspiring. And yet, not having all the answers—the very idea of *needing* to do research—can be tragic. The story Jansen tells is somehow surprising despite being painfully familiar to many.

A patient, diagnosed with a terminal illness, is invited to volunteer for an experimental treatment. Doctors say that the odds of the patient being helped are one in a million, but what the patient hears is, “So there's a chance!” Subjects know they are volunteering for a chance to produce data that may help patients in the future, but feel in their hearts that they also are volunteering for a chance to be saved.

A further twist is obvious yet disconcerting. A volunteer becomes an experimental subject, not just a patient. But serious scientific research has proper experimental controls. Accordingly, if one hundred patients are offered a chance

to be subjects in a test of an experimental treatment, some subset of the hundred will be experimental controls. They receive some kind of placebo treatment, not the treatment on which they are pinning unrealistic hopes.

If you are a principal investigator, you may find that the treatment has a beneficial clinical effect. If it does, then that means some of your subjects have enhanced prospects compared to your experimental control. What about subjects who turn out to be your controls? You hope to be proving that their grim untreated prospects are *clinically* unnecessary going forward. *Scientifically*, your justification is straightforward: namely, science identifies a treatment's efficacy by testing it. Proper controls are part of proper testing.

We all know this, but when principal investigators get subject consent, are subjects volunteering to be *controls*? Is volunteering for the "risk" of being a control the same thing as volunteering to be a control? Do we have to answer this question with a yes in order to convince ourselves that we have a right to proceed with the experiment? If that has to be good enough, does that make it good enough? How many subjects designated as controls would volunteer for the research trial if somehow they knew in advance that they would be serving as a control?

Nearly everyone sees participation in research trials as requiring free and informed subject consent. Many would say that standards for such consent go beyond what ordinary medical care requires. Today, however, Jansen sees a movement to dampen the demands of consent for research trials and in effect transcend the whole idea of a sharp distinction between research and patient care. Jansen offers a model of informed consent for this new context.

Is there a good day to die?

Suppose we are presented with two life profiles. Suppose they are virtually identical for the first eighty years. Then they diverge. One life ends a few weeks later. Someone dies in her sleep or learns that she has only weeks to live, with just enough time to gather family and friends so that she can celebrate, make amends, bid farewell, or whatever is worth doing with her final days. Then she lets go. The second life, by contrast, lasts several more years—years of terror as she loses her mind or endures physical agony or crippling disability as her body falls apart. In the abstract, just about everyone would choose the first profile over the second. Indeed, we might describe the first profile as the one where a person can die with dignity.

Real people get through real lives by observing, gathering information, and figuring out what is worth wanting as they go. Living is something we might decide that we want to do in a competent way, as is deciding that we do not. For most people in most circumstances, declining to go on living might be hard to countenance as a rational decision. Yet, Udo Schüklenk observes, there must be circumstances under which it would be rational. Moreover, we might see assessing the rationality of choosing to die as—ideally—something that individual persons have the right and responsibility to do for themselves.

What, if anything, is stopping us from having life-affirming and dignity-affirming policies pertaining to people deciding for themselves what to count as a good day to die?

Does bad law make hard cases?

There is a saying that hard cases make bad law, but it goes both ways. Flawed rules sometimes make cases hard. Greg Bognar writes about plagues, pandemics, and a contractualist alternative to cost-benefit analysis (CBA). Consider CBA's role in operations of the administrative state.¹ First, observe that CBA is more a political than a moral theory. CBA is first and foremost a theory about the challenge of real-world governance.

In the world as we know it, the rule of law is a product of negotiation and compromise, continuously evolving in ways only partly foreseen and only partly intended. We imagine that the law of the land is whatever the legislature decides. In truth, even leaving the common law aside, there is little that any legislator can simply decide. At best, legislators can hope to influence an accumulating body of law as they compromise continuously with fellow legislators who see things differently and who are honor-bound to represent different constituencies. Once a bill's final shape is resolved, legislators find themselves voting, not deciding. They may vote on short notice on a 12,000-page bill that no one in the world has read in its entirety, not even the hundreds of legislative staffers and lobbyists who each wrote a few pages of it.

If different legislative branches pass different versions of a bill, reconciliation follows. Something gets treated as having passed. Then it falls not to the legislature but to regulatory agencies (typically, agencies of the executive branch) to make sense of what just became law. Subject to judicial challenge and to the legislature's oversight, law-as-enforced will be law-as-interpreted by regulatory agencies of the executive branch. When unsure what to count as putting a bill into practice, regulators may weigh pros and cons of different options and different interpretations of their mandate. Occasionally, regulators make the weighing explicit, listing pros and cons and assigning numerical weights. That is to say, they do a CBA. Courts of justice sometimes seem to do likewise. What could go wrong? Needless to say, things can go terribly wrong. CBA is not guaranteed to pass self-inspection. However, neither are the alternatives. In the midst of a crisis such as a pandemic, our leaders will be concerned to reassure the public and avoid panic. But what if our leaders panic and imagine that managing a crisis requires stifling dissent?

Law is (literally and metaphorically) a traffic management system. Good law is an effective way of establishing and managing expectations in normal cases. It simplifies and clarifies a citizen's task of knowing what to expect, knowing what others expect, knowing how to avoid conflict, knowing how to mind one's own business, and so on, in the typical cases that structure life in a community. But

¹ See *Social Philosophy & Policy's* special issue on The Administrative State, 38, no. 1 (2021). See also David Schmidt, *Living Together: Inventing Moral Science* (New York: Oxford University Press, 2023), 159.

what about departures from the normal case? What about crises? What about cases that signal that our community will never be the same, so that what once sufficed to keep the peace and manage expectations in our society is about to change and once-adequate law is about to become obsolete?

We might worry about preparing for crisis in the manner of generals preparing to fight the last war. We prepare for a pandemic by setting aside billions of masks to decay in storage.

What is the comparably obvious problem with what we are doing today to prepare for what is coming next?

Part of our policy has to involve not expecting policy to be the answer to all questions. There are crises around the corner that we do not see coming. How do we prepare to deal with being not exactly prepared but able for all that to respond in an effective way?

Can we avoid crisis management predictably leading to bad policy? Does common law's accumulating body of judicial precedent lead to bad policy in a comparably predictable way? If not, why not? Is there any way to make it predictable that crisis managers will be on the ball? Like Bognar, Søren Holm's basic test case has to do with COVID-19, but COVID-19 is Holm's model of a larger issue. A general worry regarding health policy as it pertains to crisis management is that "hard cases make for bad law." But what is the alternative? Are some societies more resilient than others? If so, what makes a society generically more resilient?

Why is it so difficult to get things done?

Lauren Hall asks: How do we avoid creating institutions that predictably end up being guilty of structural hobbling, that is, of making it more difficult to get things done? Almost anyone who has ever worked for a large organization knows how it feels to run into bureaucratic hurdles that make it more difficult to get things done. Yet, the hurdles tend to be there for a reason: namely, hurdles make it more difficult to do damage, too.

Hall tells a story of Joseph, an elderly patient looking not for miracles but for simple comfort at the end of his life—for himself and for family members caring for him. As Hall tells the story, bureaucrats seemingly worked tirelessly to make the last year of Joseph's life as costly and as uncomfortable as possible. The small fortune that was spent was good for rent-seeking providers who knew how to game the system and cut off family access to obviously cost-effective alternative measures. The resulting torture of Joseph and his family was treated, in effect, as acceptable collateral damage.

It is not only individual human persons who commit injustice. It is corporate persons, too, including institutions that in one way or another have captured the power of government: "When [Iris Marion Young] does reference government activity, it is generally of two kinds: government's failure to act or government acting in bad faith." Hall focuses on a third category, namely, *well-meaning* government interventions that culminate in structural injustice.

Hall also recounts Mariana's story. Mariana was shut out of a career as a midwife by impenetrable layers of exclusionary licensing requirements. In both

cases, “individuals with reasonable and deeply held convictions about what should happen to their bodies or how to use their talents in society are stymied by arbitrary and seemingly senseless rules that frustrate their every attempt to navigate.” Hall offers a theory of “structural hobbling” as a contrast to the account of structural injustice offered by scholars from Young to [Maeve] McKeown.” Hall surmises that, even in the case of well-meant intervention, “theorists miss an opportunity to evaluate how interventions themselves may contribute to the burden of injustice. This injustice, like structural injustice broadly, is generally regressive; it has the most significant hobbling effect on the most vulnerable among us.” Hall supposes that “in many if not all of these cases, the community is likely worse off than it would be in the absence of structural hobbling. In Mariana’s case, the community loses her expertise and education. In Joseph’s case, the community (via Medicare) pays much more than it should for substandard care.”

On needing to start somewhere

As Carolyn Tuohy might observe, upon hearing the histories of Joseph and Mariana as told by Hall, every story has to start somewhere. Yet, every starting point has a history of its own. So, every starting point is inherently controversial. The story did not need to start there. It was a choice. A narrator chose that starting point to help illuminate the narrator’s chosen theme. Does it matter that we could dig back further into Joseph’s past in order to find inflection points where the story would have had a different ending if Joseph had chosen a different path? If we did that, what would our point be?

A history is a narrative that invites readers to interpret sequences of events as causal connections. You cannot be a serious narrator without being aware of different ways of telling a story and of different causal structures that would have been suggested by different ways of telling the story.

David Hume’s insight was that “experimental methods of reasoning” start with observations of correlation and end with inferences of causation. Francis Bacon, centuries earlier, had already seen that. Yet the world of philosophy, even in Hume’s time, continued to see science in Cartesian terms—as a search for indubitable axioms from which necessary truth could be validly deduced.

Accordingly, even in the late twentieth century, Hume was misread as a skeptic about causation and a skeptic about morality. Hume saw that observed correlation cannot entail conclusions about causation. Today, we still misread that observation as expressing skepticism about causation. In fact, Hume was attempting to apply experimental methods of reasoning to moral subjects. The skepticism he was expressing was about deduction, not causation. Hume’s explicit point was that we do not develop scientific knowledge of causation just by deducing, deriving, or proving. We have to observe, see a pattern, then jump to a conclusion that robustly predictable correlation is a sign of causation. Being scientific is no reason not to jump to a conclusion. However, having jumped, it is part of a scientific attitude to acknowledge as plain fact that we might be wrong. Seeing that scientific reasoning is not deductive and wanting readers to

understand that he was not disparaging inductive reasoning, Hume doubled down, insisting that even *ethics* is inductive rather than deductive. We gather information, then we jump to a conclusion. The jumping is not a mistake so long as it is testable, subject to correction by further observation, and so long as it is not confused with deduction.

The hell of human reasoning is that we reason in real time. We accept conclusions provisionally, as we go, then treat provisionally accepted conclusions as grounds for further and inevitably piecemeal decisions. For better and for worse, using our conclusions as grounds for further decision entrenches them and makes us less interested in disconfirming evidence. That is the nightmare of policy wonks running around grossly exaggerating the reliability of their reasons for believing what they believe, seeing themselves as compelled to pretend to know more than they actually know, and compelled to represent themselves as experts even when the issue at hand is well outside the boundaries of any field in which their expertise is genuine. Even within their range of expertise, though, they still end up telling the best story they can tell about where they think we are heading, with a narrative that reconstructs selected details from a known past so as to make the strongest possible case for predicting what they think is most likely.

Sarah Raskoff's essay is a reflection on the central importance of choosing for oneself. In the process, she illuminates another way in which a life becomes like a narrative insofar as the choices we make amount to choosing not only what we are trying to achieve, but also who we are trying to be. Decision-making not only *responds* to values but *generates* them. Values need not be exogenous inputs into decisions. Values develop and emerge from crucibles of choice. End-of-life choices, such as those discussed by Schüklenk and Hall, unmistakably have that character, but so do decisions regarding termination of pregnancy. Choosing for yourself is a path to self-discovery, to be sure, but also, and even more profoundly, self-creation.² On Raskoff's view, deciding whether to terminate a pregnancy can present itself as a hard choice. The point is not at all that we find ourselves in situations of indifference. On the contrary, the problem is so important that it is overwhelming. We are choosing the kind of person we want to have been. (Somehow, intuitively, navigating these existentialist crises well is partly a matter of having faith that we will not regret who we chose to become in those moments of existentialist crisis.)

Frances Kamm explores another aspect of the range of health policy puzzles associated with the limits of autonomy and the limits of a woman's right to choose with her discussion of how the concept of "innocent threat" influences health policy. As Kamm puts it, prioritizing the health of some people sometimes involves "restricting the freedom of or to some degree harming other people when those others present a health threat. This can be true, even when those

² Raskoff worries also about the propensity of paternalistic attempts to "nudge" to short-circuit the process of autonomous self-creation and to accordingly limit our potential to emerge from our hard choices as freestanding adults who figured out how to grow into (or to outgrow, as the case may be) their choices. *Social Philosophy & Policy* will return to this in a special issue on the topic of Economic Rationality, Political Rationality, and Nudging, 43, no. 2 (2026).

others are morally innocent threats.” Kamm recalls how Judith Thomson’s seminal 1971 essay made the problem of abortion one of the twentieth century’s most notoriously vexing philosophical puzzles.³ More generally, Thomson helped to identify a constellation of moral problems having to do with “innocent shields” and “innocent threats.” Thomson also was among the earliest theorists of related puzzles that we now call Trolley Problems. The paradigmatic Trolley Problem is a thought experiment that draws out intuitions about whether it is morally permissible for a bystander to turn a runaway trolley away from killing five people tied to the track to a sidetrack where the train will kill one person tied to the sidetrack. (Kamm’s essay recaps the crucial details, and indeed, Kamm herself has long been at the forefront of many of these issues.)

Why are Trolley Problems a problem? Not because there is anything special about the numbers one and five. Rather, there somehow is a morally special difference between more and less. There is also a question of whether anything matters other than the difference between more and less. Intuitively, other things matter, but Trolley Problems are constructed with a view to giving us cases where nothing else matters unless, perhaps, there is a difference that matters between killing and letting die. Suffice it to say, philosophical training equips us to contrive cases where differences that normally matter (in the intention, say) between killing and letting die do not matter in the particular case.

Kamm also discusses killing versus letting die. That killing is morally worse is a powerful default assumption, problematized in brilliant fashion by James Rachels, but still a default assumption for all that.⁴

Intuitively, the maxim that we should do the best we can is compellingly obvious. Less obviously, morality is not always about what to do. Sometimes, morality is about what to respect. We know there are circumstances where there is nothing for doctors to do other than do the best they can, but we also know that those circumstances are suggestive of tragedy. We know “more is better” is not necessarily true—counterexamples abound—yet it remains a default. Given an option of, say, doing five times as much good, going the other way would need explaining.

Our legal concepts draw bright lines between what is inside the boundaries of legality and what is outside. But we sense that when it comes to law, there is a compellingly nonarbitrary reason to draw bright legal lines. For example, if you fly over someone’s property at a certain altitude, will it count as trespassing?

Part of the point here is that our categories help us to sort out understandings of reality. We can divide what we observe into categories of person and non-person or into categories of beings with free will and beings without. In all cases, though, our categories are meant to simplify and thus clarify and illuminate an underlying reality. The unavoidable problem with our simplifications is that they will be, after all, simplifications, and therefore are not guaranteed to serve their intended purpose going forward. Moreover, they will be *purposive* simplifications,

³ Judith Jarvis Thomson, “A Defense of Abortion,” *Philosophy & Public Affairs* 1, no. 1 (1971): 47–66.

⁴ James Rachels proves this in his “Active and Passive Euthanasia,” *New England Journal of Medicine* 292, no. 2 (1975): 78–80.

intended to align with purposes we bring to the task of simplifying reality. We will have *reasons* for wanting to see boundaries of personhood, responsibility, free will, or the right to say no in one place rather than another.

Philosophers are trained to assume that metaphysical categories are more fundamental than legal categories. Legal categories are supposed to be derivative, nonfoundational, and accountable to a more fundamental metaphysics. Is that tenable in the case of personhood? Or do our decisions about what to regard as a person answer to more fundamental questions about who we need to regard as responsible agents in order to be able to function as a community?

We do not decide in a vacuum whether to see ourselves as having free will. Seeing ourselves that way is a prerequisite of seeing ourselves as persons. Seeing ourselves as persons, and thus seeing ourselves as having the accountability that goes with being persons, is more an ethical than a metaphysical decision. Furthermore, deciding to draw age sixteen as a line between childhood and a kind of accountability that goes with adulthood is more a legal than a metaphysical or even an ethical decision. Deciding that five is more than one is not a legal decision. Neither is deciding that saving five is more than sacrificing one. Yet, our law has decided to treat killing people as outside the purview of minding our own business, regardless of the fact that we easily can imagine a case where we could save five by killing one. The law aims to take killing off the table and has compelling reasons to do so. There is no comparable point in aiming to take “failing to save x number of people” off the table. This is a point where morality takes its cue from the law, not the other way around. “It’s against the law” is a reason grounded in legal thinking. By contrast, “it’s against a law without which we would not be able to trust neighbors (or doctors)” is moral reasoning grounded in more fundamental thinking about what sort of political and legal arrangements enable social animals to be better off living together.

Kamm also considers how philosophical problems regarding the morality of our treatment of innocent shields can bear on pandemic policy: Kamm supposes that wearing a mask “can be a duty. If someone knowingly does not do his duty, he is no longer morally innocent.” Kamm also sees “a moral asymmetry between the threatener and the threatened such that the innocent threat has a greater responsibility to avoid harming than the potential victim has to avoid being harmed.” Perhaps. In any case, this comports with our intuition that we want to avoid blaming the victim, even though we do of course want victims to take such opportunities to protect themselves as are reasonably within reach. Still, at some margin, the responsibility would seem to be shared. People who are not aware that they have a contagious disease and are putting others at risk are what Kamm calls “minimally responsible agents” insofar as they should be aware that they have been in situations where they may have picked up a contagious disease. They may be aware that they are experiencing symptoms that might be something other than a seasonal allergy. Kamm wonders whether there is a limit on counting as an innocent threat as one slides from being unaware of—to pretending not to notice—mounting evidence of being contagious.

Michael Cannon also discusses variations on Trolley cases. A point emerges from these and other discussions of such cases. Namely, we construct our stories hoping to abstract away from inconsequential details so as to focus a reader’s or a

student's attention on what we think really matters: either the difference between one and five or the difference between killing and letting die. In fact, the abstraction actually achieved by this is different because we end up focusing on what to do as if there were nothing to being a moral agent beyond deciding what to do. And yet, one main lesson implicit in these discussions is that *communities* provide agents with moral structure that *theories about what to do* cannot provide. Communities can be structures of mutual expectation that require individual agents not to imagine that morality authorizes them to go rogue and ignore what communities actually need from them. In particular, communities do not need individual moral agents to optimize in a way that is heedless of constraints implicit in what the agents around them need to be able to expect from them. Communities can generate a fabric of mutual expectations such that people learn what to expect from each other and learn what to count as staying in their lane. Communities can identify boundaries as such and in effect require individual agents to think of boundaries as settling the matter, not merely as having weight for agents to take into account.

Cannon worries about bureaucrats becoming frustrated by political gridlock and ending up deciding that they will do whatever it takes to impose their will. He worries that the imperative to get vaccinated can become an article of faith, to a point where beleaguered bureaucrats start feeling pressure to deny that there are risks. Bognar, as noted above, sees the same frustrated over-aggressive will to power in our pandemic management policies. Bureaucrats sometimes feel a need to pretend to know more than they do or pretend to have a mandate that they do not have.

Cannon finds it legitimate for government to “introduce coercion only when there is clear and convincing evidence as opposed to mere likelihood that an intervention would minimize the level of violence in society.” He says that subsidies per se are not coercive or are only as coercive as are the taxes that finance the subsidies. Cannon goes further and declares that it “introduces no additional coercion” to later impose new restrictions on those subsidies or to threaten to withdraw those subsidies. Is this true?

From narrator to visionary

What to regard as health policy in the first place has a history of being a politically fraught question. For example, Jessica Flanigan asks: Does eugenics count as health policy? Yes, she says, albeit politicized and hijacked health policy. Hijacked policy is policy that invites powerful people to hop on the bandwagon hoping to convince at least themselves that they are progressive thought-leaders. Public officials sometimes seem not to understand the situation and to be unable to adapt as the situation changes. Is part of the problem that they are seeing the situation through an ideological lens that becomes a “cocoon of confirmation bias that fosters false consciousness” and through which important new information is rendered invisible?⁵

⁵ For further discussion, see *Social Philosophy & Policy's* special issue on Ideology, 41, no. 1 (2024).

Note that while eugenics is now entertained only by extreme authoritarians, the original push for eugenics came from people who saw themselves as progressives. Today, there is nothing inherently conservative or progressive about eugenics. The will to entertain eugenic policies is today perceived simply as hubris, an astounding over-confidence in the competence and the good will of the kind of people willing to do what it takes to win political fights for that kind of life-changing power. The will to impose a vision makes mission creep inevitable.

Peter Jaworski also worries about a tendency toward an ideological understanding that can structure our perceptions in a possibly misleading way. Specifically, when we see interactions being shaped by forces of supply and demand and see a seller's willingness to serve buyers and to be influenced by what buyers are willing to pay for a given service, is that tantamount to seeing that service as having been commodified? What is at stake? Debra Satz argues, influentially and deservedly so, that commodification is a real thing, and potentially a real worry, but not a conversation-stopper.⁶ For example, Satz says, it is difficult to protect prostitutes against their bodies being commodified by taking away what some of them see as their best way to make a living under their circumstances. It really is plausible that those who would prohibit commodification are aiming (however, benightedly) to prohibit exploitation, domination, and ultimately inequality. But here, too, there is a risk that what ends up being oppressive is not the disease of commodification but the cure of prohibition. Jaworski's quarrel is specifically with regard to proposed prohibition of life-saving medical services that use donors who are in short supply, where shortages can be alleviated by allowing medical providers to compensate donors. The will to prohibit has to be educable by observing foreseen but unintended consequences of actual practices. When people are donating blood plasma to pay the rent and are making themselves useful in the process, we are queasy about people being in those circumstances and may well wonder about the safety of the supply being generated under what might feel like a degree of duress, but that sounds more like a signal to proceed with caution than a signal to stop.

James Stacey Taylor brings that conversation full circle, back to questions of informed consent raised by Jansen's lead essay. Patients should be, as much as is humanly possible, fully informed. So should experimental subjects, although (again) that is a different set of issues. As Raskoff and Schüklenk also stress, the issues concern not simply the patient's best interest *qua* patient, but also the patient's autonomy. Taylor distinguishes competence from autonomy and argues that health policy properly requires patients to give competently informed consent, not merely autonomous consent. He also argues persuasively that being autonomous does not require that all of one's relevant beliefs be true (although what he says about believing in giant raccoons will leave some readers wondering). Yet Taylor denies that concern in health policy for patient autonomy should be replaced by a concern for patient consent. Someone who lets a patient make a decision based on a false belief need not be deliberately manipulating the patient's information in an attempt to control the patient's decision. Consent can be

⁶ Debra Satz, *Why Some Things Should Not Be for Sale: The Moral Limits of Markets* (New York: Oxford University Press, 2010).

undermined to some extent even while leaving intact the conditions for autonomous choice. Whether a patient's decision is autonomous depends not only on a patient's beliefs per se, but also on the framework of relationships within which the patient's beliefs are being formed.

Competing interests. The author declares none.