A New Ethics for Learning Healthcare

1. Introduction

Bioethics is by now familiar with calls for the development of a "learning healthcare system" (LHS), where medical research and healthcare are seamlessly integrated.[1] There is not yet consensus on what an ideal LHS would look like. At minimum it would include large-scale observational research on electronic medical records. It may also include comparative effectiveness research and other interventional studies. Regardless, the goal is to make research a pervasive, routine part of care. Doing so has the potential to produce large gains in health and wellbeing, not to mention reductions in cost, as effective health interventions are promoted and ineffective and harmful ones are weeded out.

Alas, the traditional dogmas of research ethics make trouble for LHSs. On the traditional view, research and care are activities with fundamentally different aims, governed by different norms, and subject to distinct oversight structures. In particular, because research is governed *not only* by the norm of promoting the health and wellbeing of the individual patient, *but also* by the norm of promoting the wellbeing of other future patients, special measures are required to protect research participants from exploitation and abuse, including prior review of research by an independent body and the requirement of voluntary, informed consent. But these may impede the full range of research activities envisioned for an LHS. For instance, requiring a full research consent before participation may compromise the external validity of pragmatic trials

[1] L.M. Etheredge, "A Rapid-Learning Health System," *Health Affairs* 26, no.2 (2007): 107-113; E.A. Largent., S. Joffe, and F.G. Miller, "Can RESEARCH and CARE Be Ethically Integrated?," *Hastings Center Report* 41, no.4 (2011): 37-46; R.R. Faden, N.E. Kass, S.N. Goodman, P. Pronovost, S. Tunis, and T.L. Beauchamp, "An Ethics Framework for a Learning Health Care System: A Departure from Traditional Research Ethics and Clinical Ethics," *Hastings Center Report* 43, no. s1 (2013): S16-27.

investigating interventions on short, routine patient-provider interactions, e.g. testing the use of checklists by intake nurses.

Thus, the calls for the development and institution of LHSs have come alongside calls for the development of a new ethics to support them. Sometimes these calls are framed dramatically, in terms of a new foundational framework for research ethics, complete with a hitherto-unrecognized obligation on behalf of patients to participate in research,[2] or even as rejecting the conceptual distinction between research and care entirely.[3] Other times they are more relaxed, for instance, putting themselves forward as complicating conventional (but perhaps over-simple) wisdom,[4] or simply describing a sample LHS and arguing that it would be in important respects desirable.[5]

Research ethics is a large field; rethinking research ethics is a large endeavor. Here I focus on one piece of that endeavor. I consider how the context of an LHS might affect the interaction between risk and consent. This inquiry is framed in terms of three questions. First: in an LHS could it be permissible to make some research participation compulsory? Second: by what mechanism could participation be compelled? And third: what, if any, risk limit is appropriately applied to research that is so compelled?

In existing discussion, proponents of LHSs have allowed that some research might be conducted without full voluntary consent, but have focused their attention exclusively on research which is either close to, or within, the boundary of minimal risk.[6] They have, e.g., discussed the administration of interviews,[7] before-and-after comparisons of quality

---

[2] Faden et al., "An Ethics Framework"
[3] N.E. Kass, R.R. Faden, S.N. Goodman, P. Pronovost, S. Tunis, and T.L. Beauchamp, "The Research-Treatment Distinction: A Problematic Approach for Determining Which Activities Should Have Ethical Oversight," *Hastings Center Report* 43, no. s1 (2013): S4-15.
[4] L. Gelinas, A. Wertheimer, and F.G. Miller, "When and Why Is Research without Consent Permissible?," *Hastings Center Report* 46, no. 2 (2016): 35-43.
[5] Largent et al., "Can RESEARCH and CARE?"
[6] Largent et al., "Can RESEARCH and CARE?," 39; Faden et al., "An Ethics Framework," S25.
[7] Faden et al., "An Ethics Framework," S23.

improvement interventions with few conceivable downsides,[8] and/or the randomization of patients between treatments for which there is as of yet no clinical evidence supporting either over the other.[9] And even in such cases, they tend to take an only-insofar-as-absolutely-necessary attitude, e.g. allowing opt outs and appeals,[10] stipulating that consent can *only* be waived when infeasible to obtain,[11] requiring both broad and specific forms of notification.[12] It's worth noting that, in this respect at least, the new systems they envision do not depart that significantly from the status quo. Minimal risk research is *already* eligible for waivers of informed consent under the US regulations, and the conditions under such waivers are made available are remarkably similar to those described by proponents of LHSs.[13]

It's natural to wonder whether it's possible to go further: in the brave new research ethics, is it ever permissible to enact compulsory research that crosses *beyond* the boundary of minimal risk? I answer "yes." In making my case, I begin by reviewing two prominent arguments in favor of adopting LHSs. First, I show that although those authors have limited their conclusions about compulsory research to that which is no more than minimal risk, the reasons they give do not actually appear to support such a limit. Then, I go on to develop one of those proposals more systematically. The result is an argument that an LHS could legitimately make at least some forms of more-than-minimal risk research compulsory, including both a recipe for constructing cases where this is so and a proposal for how such research could be implemented.

In sum, the final answer to my three questions ends up being: "yes, participation could be made compulsory; if so, the threat of withholding care is a legitimate mechanism for that

---

[8] Gelinas et al., "When and Why?," 38.
[9] Largent et al., "Can RESEARCH and CARE?," 39
[10] Largent et al., "Can RESEARCH and CARE?," 41.
[11] Gelinas et al., "When and Why?," 35-36.
[12] Faden et al., "An Ethics Framework," S24.
[13] It is worth noting that Gelinas et al. take the Common Rule's already-existing *absence* of a universal requirement for consent as their starting point. Gelinas et al., "When and Why?"

compulsion; and it is legitimate to employ such compulsion whenever the burdens of participation, relative to the gains of research, stand in no worse a relation than do those that patients must already accept in the course of cost-benefit rationing as employed in ordinary clinical care."

2. Faden et al.: Reciprocity and Common Purpose

Faden et al. have used the occasion of a special issue of the Hastings Center Report to put forward a new basic framework for research ethics, one designed to accommodate the integration of research and care.[14] That framework contains seven obligations. The first four concern the protection of the patient/subject and their relationship with their clinician and the fifth enjoins the system to address health inequalities. For our purposes, it is the sixth and seventh that are important: they command that learning activities be conducted and that patients participate in research.

In discussing the grounds of this obligation on patients, Faden et al. cite both a norm of reciprocity and a norm of common purpose. The former directs us to promote conditions that are in the interest of all, particularly when that can only be achieved by the participation of all, and the latter directs us to return the benefits we receive from others proportionally. Because medical progress is in the interest of all, and because all benefit from those who participated before, both support an obligation to participate.

What does this obligation amount to in practice? Faden et al. emphasize that balancing the obligations they posit against one another requires judgment. Nonetheless, they do offer concrete illustrations. In particular, they indicate that the obligation for patients to participate will be limited by considerations of risk: patients are not obligated to participate in high risk trials of

[14] Faden et al., "An Ethics Framework"

new investigational devices, but may be obligated to participate low risk research activities, such as entry in registries, having their de-identified records reviewed, giving interviews with staff.[15]

These examples are *so* low risk, though, as to negligible: participating in an interview, for instance, is a paradigmatic "minimal risk" form of research. But why think that these putative norms of reciprocity and common purpose hold only up to the extent that the research in question is minimal risk? The norm of common purpose certainly doesn't explain why examples should be selected only from this class—after all, it is not as if *only* minimal risk research contributes to our common purposes. Higher risk research does too.

It is hard to see how the norm of reciprocity could explain a minimal risk limit either. The gains we receive from past medical research are enormous. For instance: by one report, in the 18th century smallpox claimed the lives of approximately 2,000 out of every million living Swedes annually.[16] The eradication of that *one* disease is already a substantial benefit to each and every human, but even within the category of vaccines, we must also count measles, rubella, diphtheria, polio, etc. And vaccines are just one area where generalizable medical knowledge has improved human health.

Compare these benefits against the risks that fall within a minimal risk threshold. Those risks are defined in US regulations as those no greater than the ones encountered in daily life or routine physical and psychological examinations. One comparator used to establish "risks of daily life" is the risk of travelling by car, given that car trips are a part of daily life. To get a feel for this standard, trying focusing on a particularly risky kind of car trip: an adolescent driver, on a rural highway, driving in rainy conditions might have something like a 4 per million chance of

---

[15] Faden et al., "An Ethics Framework," S23, S25. These examples are not fully helpful. Review of de-identified medical records is not "human subjects research" under current regulations, and so already does not require consent or notification. J. Menikoff, "The Unbearable Rightness of Being in Clinical Trials" *Hastings Center Report* 43, no. s1 (2013): S30-31.

[16] Internet Archive, "A concise history of small-pox and vaccination in Europe," August 2019 at https://archive.org/details/b21357262/page/n26. As reported there, the statistic is attributed to a report by the Swedish Statistical Bureau to the German Imperial Health Office.

death.[17] As a very risky trip, that presents a possible upper limit. Although there is disagreement, few would ever countenance as "minimal risk" a mortality risk of greater than 10 per million.[18]

It is tricky how to compare single-instance risks (risk per drive) against durable protection from smallpox. Perhaps the best comparison, then. is annual risk to annual risk: in 2015 Sweden had an annual traffic fatality rate of 28 per million, as again the above-quoted 2,000 per million fatality rate it once experienced from smallpox.[19] The difference for just that one disease is already orders of magnitude. In any case: because Faden et al. do not give any explanation of how they understand the notion of proportionality that's involved in proper reciprocity, it is difficult to know the exact comparison to make in order to *demonstrate* that figures like these are out of proportion with one another. Nonetheless, I see no obvious way of understanding proportionality such that accepting *only* minimal risks ends up being a proportional way of reciprocating the full gains of human participation in medical science.

Perhaps Faden et al. would respond that the minimal risk threshold derives not from anything about the obligation to participate on its own, and so it need not be explained by the norms that justify that obligation. Rather, it arises only once that single obligation is balanced against others in the framework, including those which generally protect patient interests. But this is unsatisfying. First, even accepting that these obligations are to be balanced, there is no explanation for why the balance is struck in this particular way. Second, it is independently implausible that this represents a genuine balance.

---

[17] D. Wendler, L. Belsky, K.M. Thompson, and E.J. Emanuel, "Quantifying the Federal Minimal Risk Standard: Implications for Pediatric Research Without a Prospect of Direct Benefit," *Journal of the American Medical Association* 294, no.7 (2005): 826-32.

[18] Shah et al. found 93% of surveyed chairs of pediatric IRBs were unwilling to categorize a 1 in 100,000 chance of death as "minimal risk." S. Shah, A. Wittle, B. Wilfond, G. Gensler, and D. Wendler, "How Do Institutional Review Boards Apply the Federal Risk and Benefit Standards for Pediatric Research?" *Journal of the American Medical Association* 291, no. 4 (2004): 467-82.

[19] WHO, "Violence and Injury Prevention," August 2019 at https://www.who.int/violence_injury_prevention/road_safety_status/2015/GSRRS2015_data/en/

Consider that just about *any* time a norm of reciprocity or common purpose applies, there will also be some countervailing considerations with respect to "protecting" the obligated person's interests against infringement. The point of positing an obligation in the first place is to explain why those interests should sometimes be infringed, contrary to what a person might otherwise find most convenient. Simply citing the existence of such interests doesn't say much. If there is a special reason why the medical context is different from any other in this respect, we have yet to hear it.

3. Largent et al.: The Veil of Ignorance

Largent et al.'s proposal for how to desirably structure an LHS has many elements, but again we can focus on one: the "prescribed trial."[20] When an LHS offers a prescribed trial, it stipulates that participation is mandatory if an enrollee is to have their medical care reimbursed at normal rates. So receiving full compensation may require, for instance, accepting that one will be randomized between standard therapies as part of a comparative assessment study. Should an enrollee refuse, they can still access standard care but must pay an additional surcharge. What trials will be prescribed? Largent et al. also endorse a minimal risk limit. Anything riskier must be instead "invited," a distinct category without sanctions attached to refusal.

What justifies the system in deciding to prescribe trials? Largent et al. argue first that the burden of pseudo-mandatory minimal risk research is negligible and that the potential gains are large. This argument is sensible, but it is unclear how much it ultimately illuminates—in particular, it does not give us a theoretical handle on what happens once we start considering ever-so-slightly less negligible risks and as-large, larger, or slightly-less large gains. But that is the question of interest to the present paper: how far ought we to be willing to go?

[20] They refer to it as a RIHS or "research integrated health system," but this terminological difference is immaterial. Largent et al. "Can RESEARCH and CARE?"

To make progress on that question, it is more useful to consider another separate, more theoretically robust justification Largent et al also offer: they argue that a system which prescribes trials would be preferred if one were to consider the question of which health system one would like to join, abstractly and from the outside. In articulating these grounds, they cite John Rawls' famous veil of ignorance device.

But what about the justification provided by considering things from the outside suggests that prescribed trials ought to be limited to no more than minimal risk? Recall that Rawls famously thought that his veil of ignorance would force agents to select a form of social organization aimed at maximizing the fortunes of the worst off. It is not obvious how this bears on health. But suppose that we were to use the distribution of ailments as our measure of advantage: the healthy are maximally well off, and the sick, to varying degrees, not so. But then a distributive principle directing that we maximize the fortunes of the worst off, regardless of cost to the better off, would seem to support routinely prescribing high risk trials on healthy volunteers. For a deliberately extreme example, consider requiring healthy volunteers to undergo Ebola infection in carefully monitored "challenge" trials. So long as the trial still leaves participants better off than contracting Ebola outside the context of a trial, the healthy "volunteers" would not be among the worst off, and insofar as the burdens they accepted actually acted to aid the worst off the imposition of those burdens would be required by justice. Those risks are not minimal.

To be clear, the above is not intended as a rigorous discussion of Rawlsian health justice, which is beyond our scope (more on this shortly).[21] It aims only to suggest that there is a serious lacuna in Largent et al.'s argument, insofar as it is unclear what about their proposed justification for prescribing trials *at all* explains why they should be *only minimal risk*.

[21] Not least because the most prominent Rawlsian account of health justice does not actually locate the distribution of health under the difference principle, but instead under fair equality of opportunity. N. Daniels, *Just Health: Meeting Health Needs Fairly* (New York: Cambridge University Press, 2007)

3. The Political Approach

Largent et al. suggest that we can defend an LHS with reasoning relevantly similar to Rawls'. I believe they are on the right track, although nothing so strong as Rawls is required. Even very general and relatively noncommittal arguments made using the tools of political philosophy can show that an LHS can not only be just, but that it can justly compel participation in more than minimal risk research.

In making this argument, I bracket questions about the "system" that will be "learning." For instance, I do not assume an LHS must take the form of a single state-run national system, like the British National Health Service. It may be one or many entities performing research, and they may be public or private. Rather, my assumption is that *if* there are many, potentially private entities, then they are nonetheless effectively regulated such that the knowledge gained from their learning activities translates into fairly distributed future improvements in care, rather than being captured solely as private profits. The best strategy for accomplishing that is beyond the scope of this paper. Instead, I just ask: *given* that there is more-than-minimal risk research on some that, if conducted, is expected to lead to health gains for others, under what conditions should providers be authorized to make participation mandatory?

In answering, I operate within a broad, ecumenical framework. It says the state should generally pursue policies that have *good distributive effects*, in terms of improving the amount of some relevantly defined goods and the degree to which the distribution of those goods across people conforms to a desirable shape. It also says that in pursuing that aim it should take itself to be constrained by *individual rights*. That is, the state is barred from doing at least some things even if they would otherwise promote its distributive ends.

In the rest of this section, I consider the possibility of enabling compulsory research from the perspective of its distributive effects. Sometimes, in thinking about distribution, one must

specify precisely what is distributed. But here we need not decide whether those benefits and burdens consist in e.g. welfare, resources, or capabilities. Similarly, sometimes one must specify the desired shape of the distribution, but here we also need not decide whether e.g. distributive shape doesn't intrinsically matter at all, as in utilitarian analysis, distributive concerns ought to somewhat favor the less advantaged, as in prioritarian analysis, or distributive concerns require decisive weighting of benefits to the least well off, as Rawls famously argued.

Regardless of what the fundamental currency of distribution consists in, any plausible view will imply that the morbidities and mortality associated with poor health status count *at least instrumentally* as bads to be prevented.[22] So think about distribution of health while holding all else equal. Consider the prospect earlier described of healthy volunteers being required to accept Ebola infection under carefully monitored conditions. The morbidity and mortality rates for Ebola virus disease when treated with aggressive supportive care are substantially lower than the morbidity and mortality rates when such care is delayed or unavailable.[23] Hence, if performing medical research which involved inflicting one case of Ebola *with aggressive care options available* could yield health knowledge that went on to prevent another case of Ebola in a context where those measures would not be available, that tradeoff would (all else held equal) be distributively desirable. It would be desirable on brute utilitarian grounds, minimizing the total amount of burden, but also on prioritarian and minimaxing grounds, as the most burdened person would be less burdened. Regardless of our more specific view, such research looks distributively desirable.

---

[22] The connection is obvious for welfare and some capabilities accounts, less obvious for some other capabilities and resource accounts. But to the extent that those latter accounts really cannot secure even an instrumental role for paradigm instances of good and bad health, that itself discredits them.

[23] T.M. Uyeki, A.K. Mehta, R.T. Davey, A.M. Liddle, T. Wolf, P. Vetter, S. Schmiedel, et. al 2016, "Clinical Management of Ebola Virus Disease in the United States and Europe," *New England Journal of Medicine* 374, no. 7 (2016): 636-46; L. Aleksandra, W.A. Fischer, T.M. Uyeki, T.E. Fletcher, N.K.J. Adhikari, G. Portella, F. Lamontagne, et al., "Ebola Virus Disease and Critical Illness," *Critical Care (London, England)* 20, no. 1 (2016): 217.

Of course, one does not know in advance that a particular research program *will* yield a given result, and so one cannot guarantee that inflicting a case of Ebola here could prevent a case of Ebola there. One would have to discount according to the likelihood of research failure. At the same time, those research programs that *do* succeed affect huge numbers of people. Even just accelerating an already-inevitable breakthrough may translate into millions of lives positively affected. While estimating the effects of any given research presents a difficult technical challenge, there is no apparent reason to believe that it will never come out looking desirable.[24]

What the Ebola case dramatizes is that the size of a discrete benefit or burden *in isolation* isn't the grounds on which it is distributively evaluated. Even very large burdens can be justified by the prevention of yet larger ones. As a result, distributive concerns do not appear to be of the right form to recover a minimal risk limit, as whether a burden is greater or lesser than minimal risk *is* determined in isolation. Furthermore, it is not just that, conceptually, burdens *can* be linked, and that large burdens can sometimes prevent yet larger ones: it's that illness at least sometimes appears to have those characteristics, as although being forced into research can be burdensome, being very sick is one of the most burdensome things that there is.

As a result, it strikes me that if there is to be an objection to conscription for more than minimal risk medical research, it does not stem from distributive considerations. Still, that leaves it open that such conscription might nonetheless be ruled out by those individual rights that we earlier allowed constrain the state in how it can pursue desirable distributions. Medical research

[24] One might object that the foregoing discussion has critically overlooked the difference between identifiable and statistical lives, or, alternately, between risks considered *ex ante* and outcomes *ex post*. I reply: if that were the *only* problem, then we could resolve it just by ensuring our inability to know who was to be conscripted. Imagine that every hospital in the country was mailed bushels of unlabeled vials. Almost all are saline and inert; a few will cause Ebola infection. The *ex ante* risk from being injected with a vial is both equally distributed and miniscule in magnitude. The people who fell ill could then be transferred to the ICU and treated during the course of research.

often involves trespass on the body, and many believe in some morally important form of individual sovereignty over the body. Perhaps that explains why authorizing compulsory risky research is illegitimate: even if carrying it out would have good distributive effects, doing so is ruled out by prior considerations of individual rights.

I will, going forward, make one major concession to such concerns. It is reflected in what I will mean by "compulsory" research. Although I am willing to entertain the thought that it could be justifiable to compel research similarly to military conscription—a lottery subsequently enforced by criminal law[25]—I am simply not going to talk about that here. Instead, I focus my attention on a version of what Largent et al. referred to as "prescribed trials." When I discuss an LHS making more than minimal risk research "compulsory," I refer to that LHS making participation in research a mandatory condition on their providing care. This, though, is consistent with allowing that some may opt out rather than comply. Some may go untreated; others may access other sources or engage in medical tourism. This, in turn, does preserve a fundamental measure of bodily autonomy. The body will only be invaded when individuals present themselves for care, are adequately informed, and consent to the relevant procedures. No one will be held down.[26]

Suppose, then, the state is considering authorizing LHSs to treat participation in research as a prerequisite of receiving care, although allowing that patients must still be adequately informed and left free to opt out of accessing the system entirely. Is there anything

[25] See for instance C. Fabre, *Whose Body is it Anyway?* (Oxford: Oxford University Press, 2006), which endorses, inter alia, a non-vital organ draft.

[26] There is a tension here between bodily rights and fairness, insofar as medical tourism is not equally accessible to all. Holding care hostage may exploit those with fewer options. This tension deserves scrutiny. However, I maintain that although it's fairly clear in what manner a system which allows voluntary opting-out may be less fair than one using conscription, it is not at all obvious why it would be less fair *than the status quo*. Hence, this concern strikes me as part of an intramural dispute between competing visions for a LHS, not an objection to the project of instituting one.

about our *individual rights* that further constrains how this form of compulsion can be legitimately deployed against us?[27]

4. Individual Rights and Treatment Menus

In this section I argue that the concession I just made—that the compulsion in "compulsory research" is to be limited only to the threat of withholding care—is already enough to make it clear that compelling submission to medical procedures in this manner does not *generally* violate individual rights. A cursory examination of ordinary clinical care shows that such compulsion is routine; there is no remotely plausible theory on which all of it is rights-violating; hence, at least some of it is not.

To elaborate: contra the traditional gloss on the distinction between research and clinical care which I gave in the introduction, in clinical care health providers routinely *do* make decisions that burden some and benefit others—that is, decisions which are often enough not in the best interests of the individual patient before them.[28] For instance, physicians may decide to use narrow spectrum antibiotics, despite their reduced individual effectiveness, because doing so reduces the likelihood of future patients being affected by antibiotic resistant bacteria. A clinic may schedule a patient to receive a procedure from a new physician instead of from a more experienced practitioner, for purposes of balancing staff schedules and training future talent. A

---

[27] This question is similar to the one (Gelinas et al.) use to frame their discussion, though they are deliberately quietistic about what patients' rights actually are.

[28] D. Wendler, "Are physicians always obligated to act in the patient's best interests?," *Journal of Medical Ethics* 36, no. 2 (2010): 66-70; D. Wendler and R. Johnson, "When clinical care is like research: the need for review and consent," *Theoretical Medicine and Bioethics* 37, no. 3 (2016): 193-209. Furthermore, the fact that these practices are not only tolerable but also tolerated is important, as some authors who have questioned the normative significance of the body per se--including Wertheimer, and also Eyal--ultimately emphasize the occasional need to defer to what the public is willing to tolerate. But the public does tolerate these analogous clinical practices. A. Wertheimer, "(Why) should we require consent to participation in research?," *Journal of Law and the Biosciences* 1, no. 2 (2014): 137-82; N. Eyal, "Informed consent to participation in interventional studies: second-order in a different sense," *Journal of Law and the Biosciences* 2, no. 1 (2015): 123-28.

hospital may decide to reform rather than to shutter a troubled surgical program, despite knowing that for the foreseeable future it will continue providing lower quality care.

These choices can increase the discomfort of the procedures—consider the trainee who struggles to find the vein. They can also include more serious risks, e.g. poorer post-surgical outcomes. Yet these impositions are nonetheless routine, and, in many cases, viewed as unproblematic. But how can such decisions, often made by private actors and to the disadvantage of the patient, manage to avoid violating the patient's rights?

The patient retains their rights insofar as they must consent to intake and to any clinical treatments they are to receive. Furthermore, they have informational rights with respect to what providers must tell them about the procedures they are to receive. These include a right to information that a reasonable person would take to be material to their decision of whether to consent to care. If a person initially declines care, or withdraws their consent at any time, they are free to leave and seek care elsewhere. They are free to decline to seek care entirely.

This amounts to being guaranteed that before anyone invades one's body, one must be given a clear menu of potential treatments and offered a free choice of whether to accept one or none of them. But suppose this is done: what must be explained, is explained, and the patient's consent is solicited and respected. Does the patient also have a further right to not only choose from the menu presented, but to reject the menu itself? To demand that it be replaced with another that suits their interests better? It is hard to see how they could. If patients could not only choose whether or not to accept treatment under a particular arrangement—say, a given policy for rationing antibiotics—but could, as a matter of right, insist that some other arrangement be used instead, the result would be an irresolvable impasse as soon as the first time it occurred to two different patients to demand incompatible arrangements, e.g. the first time each thought to demand to be the *sole* recipient of broad spectrum antibiotics.

So although it is plausible enough that individual rights require both disclosure and the solicitation and honoring of consent, it is not at all plausible that individual rights extend to

individual control over the terms under which the to-be-consented-to services are offered. No one's input is solicited on whether they would prefer the hospital to use a different antibiotic policy in their individual case, and no one's rights are violated by that fact. But the form of compulsory research this article investigates just *is* a modification of the terms under which a service is offered: the proposal is that the terms under which service is offered could be made to include research.

Still, even if it is clear that there is no general individual right to veto the terms under which medical care is offered, neither does that establish that anything goes. Perhaps *the reason* we all agree that patients have no right to veto antibiotic ratioing is that antibiotic rationing is independently reasonable. By contrast, one might still think that it would be unacceptable for that same hospital to maintain a surgical program which their internal data showed to be catastrophically bad, and that it would be unacceptable *even if* they clearly and fully explained how bad it is and *even if* patients' consent to surgery were freely solicited and honored. Indeed, such worries are particularly forceful when, as a practical matter, the ability of patients to decline treatment and "shop" for care elsewhere is sharply limited.

What we would like to complete our argument is an explanation of what separates reasonable treatment menus from unreasonable ones. A full account is beyond the scope of this paper. Nonetheless, here is one condition that seems unassailable: it is always reasonable for a health provider to restrict its menu of offered treatments to exclude interventions that are cost-ineffective. If patients want cost-ineffective treatment, they can seek it out; as noted, they can always decline treatment and go elsewhere. But, as I argue in the next section, accepting the claim that it is always acceptable restrict a treatment menu to only cost-effective choices is already enough to establish that a LHS may legitimately compel more than minimal risk research.

5. Prescribed Trials as Cost Effective Healthcare

Performing cost-effectiveness assessments of healthcare interventions involves quantifying their costs, quantifying their benefits, and then assessing the ratio between them. This conveys how much a provider of that intervention is paying in costs per "unit" in benefits, with the aim being to understand how to pay as little as possible while getting as much as possible. Beyond that broad formulation, there are a number of controversies with respect to the proper form and significance of cost-effectiveness analysis. For instance, there are controversies with respect to how the costs and benefits should be understood—e.g., whether narrowly related to health or incorporating broader personal and social outcomes—as well as what the obtaining of a "better" or "worse" cost effectiveness ratio tells us about different interventions—e.g., whether in case of rare diseases we should be willing to fund interventions that are less cost effective, or how to handle the budget impact of interventions that are extremely expensive in absolute terms yet also extremely effective.

Despite the plentiful disputes over how best to understand and apply it, it seems clear that there is no alternative to performing cost-effectiveness analyses. Health systems *must* decide how to allocate finite resources, and it would be absurd to do so without reference to how effective its various options were. But, I argue, once we have already accepted that it is legitimate for a health system to ration a patient to a treatment that is worse for them than another available treatment option—because it is more cost effective, once we take full account of price, side effects, benefit profile, etc.—then there is no principled reason to suppose that it could not, using the same analysis, also ration a patient to participation in a research protocol that is worse for them than some otherwise possible course of clinical care. If no one's rights are compromised by the one, neither are they by the other.

We can illustrate this with an extended example. The numbers are chosen for tractability, not realism; they will be unrealistic. The point is to illustrate a form of reasoning.

So: suppose Drug A is a name brand drug treating a chronic, unpleasant, but not dangerous condition. Drug A costs $100 per month, and the relief of the underlying condition yields 10 “units of benefit” per month, however that has been specified. Drug B is a cheaper generic treating the same condition about equally well, also yielding 10 units of benefit, but additionally has 1 unit of unpleasant side effects, for a net of 9 units of benefit overall. However, it only costs $50. The *incremental* benefit of prescribing A over B would be a gain in 1 unit at a cost of $50 dollars. Suppose our health system justifiably values units only at $30. Then, absent some further story, it would not be willing to offer A. Patients would instead be offered the slightly-inferior but much-cheaper B. This is the first step: because of the incremental increase in cost-effectiveness, the health system is justified in *only* offering B.

Now consider the following situation. Suppose that Drug A is not yet specifically approved for use in pregnant women, but nonetheless it is widely being used off-label and there is robust anecdotal evidence supporting its safety and efficacy. Consequently, there is interest in running a pharmacokinetic study of A in pregnant women with the primary objective of assessing optimal dosing and the secondary objective of ruling out potential teratogenicity. Suppose that this study would last one month and only enroll one participant. The study intends to enroll a participant who already would have received A off-label, but will now instead receive it in the context of the trial. The research component will consist in the analysis of the participant’s medical records, the collection of a few blood draws, and one lumbar puncture. Suppose that the unpleasantness—and associated health decrement—of all that is valued at 1 unit. Drug A-as-offered-in-this-trial then yields its 10 units of benefit less the 1 unit of additional research-related burden. Suppose that nothing sets out this participant as unusual with respect to how they would experience these benefits and burdens: the assignment of magnitudes is accurate not just to the general population, but also to the participant personally. Then from the

perspective of the participant, getting Drug-A-as-offered-in-this-trial for a month is indifferent with getting Drug B for that same month.[29]

Now suppose that the expected health gain from running the trial is 27 units. That is: the estimated health gains from completing the trial, translated into the same units we were using to measure clinical benefits, is "27." Now, at this point, one might object that no one could actually construct such an estimate, due to the inherent uncertainty and nonlinearity of clinical research. To that I respond: although attempting such estimates is an enormously technically challenging task, it is one we already have to live with, and that we already do sometimes carry out in a rough and ready way. That is, research funding agencies charged with promoting health already do sometimes estimate how promising various routes of scientific research are and direct their funds accordingly. So suppose that they have done so, and that the health expectation they associate with completing this trial is 27 units.

Finally, suppose that trial expenses will be an additional $100, such that administering A-in-the-trial now ends up costing $200. It costs $200, but produces *36* units of health benefit (both the 9 units to the individual and the 27 units of trial-related gains). The cost-effectiveness ratio of A-in-trial is then identical to the ratio for B, and the incremental gain in cost effectiveness of rationing people to A-in-the-trial is equivalent to the incremental gain in cost effectiveness of rationing people to B. But given that we said above that this gain was sufficient to justify overriding patient preferences between two different drugs (A & B), I see no reason not to conclude that it is also sufficient to justify overriding the patient's equally strong preferences between taking the same drug either inside or outside of a trial (A & A-in-the-trial). Hence, I conclude that it is legitimate for a LHS to prescribe participation in the trial just described.

[29] For a dramatic case where this is (presumably) violated, see Menikoff, "The Unbearable Rightness." Even if the population level health benefit of breast-sparing versus radical mastectomy is the same, as measured by e.g. QALYs, this is compatible with every individual valuing one or the other much more highly.

The above is just one example. The form of reasoning, though, is general: begin by imagining that any given case of research participation was itself a fictional novel therapy. Consider the burdens of participation as if negative side effects of that therapy, that, so considered, make it less cost effective; but also consider the expected health gains from research as enhancers of its cost effectiveness. If the resulting balance is sufficiently good that a health system would be justified in rationing patients to the fictional therapy, so too is it justified in prescribing participation in the corresponding research trial.

Granted, the case was extremely simplified. For instance, the degree of uncertainty surrounding the potential harms and benefits associated with genuinely novel interventional research will typically be much higher than those associated with either a well-understood generic or a few blood draws and an LP. That complication suggests that the above reasoning would be more easily reproduced when considering comparative effectiveness research than research into novel interventions.

Another natural concern is that the trial imagined above burdens the one person compelled to participate uniquely, whereas the burdens involved in cost-effectiveness rationing of name brand drugs may be thought fall more equally upon all. More realistic cases will of course typically involve more than one person—but so long as they do not involve the entire patient population, they will distribute burdens unevenly; even if they did, there will sometimes be distinctions between active arms and controls, and so on. So it may be thought that the compulsory research trials will generally not parallel compulsory cost-effectiveness so closely after all. But it is not true that the burdens of cost-effectiveness rationing are always equally shared. It may be that a more effective or less noxious drug is used specifically for a population at higher risk of reaction, but others are required to receive (or: start with) the generic. Here rationing is uneven due to uneven effect. Similarly, cost may be uneven, as a system may rationally decide to provide an intervention only in a city clinic, where it is feasibly affordable, even knowing that this will make access prohibitive for far-flung rural members. Finally, we are

already familiar with a case where there is simply a hard limit on a non-divisible set of benefits, and this requires distributing them unevenly, namely, organ donation. In any case, the fact that the risks of compulsory research may of necessity fall unevenly requires attention, but not in a way that undermines the fundamental idea of justifying compulsory research by drawing appropriate parallels to cost-effectiveness rationing.

Exactly where and how those parallels will be valid depends on what position one takes on the rich disputes that already exist within the cost effectiveness literature, disputes which it is generally beyond the scope of this paper to address. And, depending on the stance one takes within those disputes, one may indeed end up deriving some limits on the kinds and levels of research risk that could be justified as above—for instance, if one thinks that it is legitimate to ration people to more cost-effective treatments only when those treatments are *slightly*, rather than *significantly*, inferior to other options. Regardless, the important point for our purposes is that even if cost effectiveness rationing does have such limits, it is quite implausible to suppose that they are so strict as to include a minimal risk limit. No one would (or should) think that a safer formulation of a drug that leads to 10 per million fewer fatal allergic reactions is worth paying *literally any price* for. At some point, considerations of cost will override that reduction in risk. But as we might recall, 10 per million lies above the upper limit for most assessments of "minimal risk." We can conclude, then, that whatever else is true of them, appropriate parallels to cost-benefit reasoning will not be constrained to justifying only research that remains under the threshold of minimal risk.

6. Possibilities for Implementation

Consider an objection. It goes: traditional cost-benefit analysis balances the contemporaneous health benefits of an intervention against the cost of providing it. In analyzing prescribed research, I added in the benefits to future patients on the benefits side and the

burdens to current patients on the cost side. These costs and benefits are no longer synchronous, nor do they necessarily occur within the same system. But, one might think, insofar as cost-effectiveness rationing is justified, it is justified because the entity performing the rationing can *itself* use the savings to provide health benefits to other patients *now.*

I am dubious that the normative claim about rationing here is correct: notice that the already-mentioned cases of restricting antibiotic access and physician training also impose burdens on patients in the here and now in exchange for future benefits of uncertain magnitude, benefits which are also not guaranteed to manifest in the same system (as patients, resistant bacteria, and well-trained physicians all travel). Even so, I am willing to allow this objection for the sake of argument: allow that health providers can *only* legitimately use cost-benefit rationing when the savings they realize can be turned into immediate health gains in the here and now.

But then consider the following arrangement. Suppose that we start with a health funding agency, one which considers a large portfolio of research proposals. Suppose that they estimate, for each, the health gains associated with performing the research and how much those gains are "worth" under their institutional view of health promotion. Suppose they then use those estimates to create a willingness-to-pay portfolio. A hospital, insurance network, or other provider could then look through, choose to run a trial, and receive subsidies in the amounts reported. In making that decision, they could consider the services provided in the trial *as if* they were a new drug whose cost and effects they were evaluating (which they sometimes may be, and sometimes may not). Perhaps the trial introduces research-related risks and burdens. Those would be decrements to the benefit, considered as-if side effects. But it would also come with the subsidy. That would be a decrement to the cost, considered as-if cheaper to procure. If, after making such adjustments, the incremental cost effectiveness ratio compared against available treatments were to be decisively favorable, then the health system could accept the funding offer and decide to prescribe participation. Patients at that system would then have the

choice of whether to access service through the prescribed trial, or to attempt to find healthcare through other channels.

For all three parties to this situation, their decisions under the arrangement I just described look awfully like the ones they make under the status quo. Funders evaluate research proposals and decide how best to allocate their money. Health systems make assessments of the clinical effects of interventions and then, comparing them to their cost, decide what the best rationing decisions are. Patients then must decide whether to accept rationing or to take themselves elsewhere, just as they already do with respect to e.g. accepting a generic or finding alternate access to a branded drug.

It may be that maintaining such a separation of roles is not actually the most efficient way to fund research or provide care. Nonetheless, the point is to provide a proof of concept. In the system as described, everyone's decision making and the benefits and burdens they face look substantially similar to how they already do—with the difference being that here their actions are brought into an alignment that systematically produces health gains from research.

Throughout this paper I have advocated what will strike many as a dramatic proposal—the compulsion of more than minimal risk research. But the proposal above illustrates that such compulsion can be made surprisingly consistent with the sorts of interactions with and within the medical system to which we are all already fairly accustomed. When it comes to the chances that such reform could ever plausibly be adopted, then, I say: people already accept, to a limited but growing degree, that cost-effectiveness rationing of clinical care is desirable and just. If prescribed trials were done carefully and well, perhaps people would come to accept them too.

* * *