Compulsory Research in Learning Healthcare: Against a Minimal Risk Limit

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] But there is not yet consensus on what this integration should entail. At minimum, an LHS would facilitate large-scale observational research on electronic medical records. It may also facilitate 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 both large gains in health and large reductions in cost as effective interventions are promoted and ineffective ones are eliminated.

Despite these potential benefits, it is not clear how research and care can be integrated consistently with traditional research ethics. Research and care are traditionally viewed as activities with fundamentally different aims, governed by different norms, and subject to distinct oversight structures. Research is governed *not only* by the norm of promoting the health and wellbeing of participants, *but also* by the norm of promoting the wellbeing of others, which can sometimes require imposing additional burdens on participants. Consequently, special measures are required to protect research participants from exploitation and abuse. These include the prior review of research by an independent body and, in most cases, a more involved informed consent process than is typical of clinical care. But such requirements 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 investigating interventions on short, routine patient-provider interactions, e.g. testing the use of checklists by intake nurses.

Hence, bioethicists have begun exploring new approaches to research ethics which may be more compatible with a robust role for learning healthcare. Sometimes these views are framed as radical departures from tradition, positing a new foundational framework for research ethics complete with a hitherto-unrecognized obligation on behalf of patients to participate in research,[2] or even rejecting the conceptual distinction between research and care entirely.[3] Other times they are framed more modestly, for instance putting themselves forward as complicating conventional wisdom,[4] or simply describing a sample LHS and arguing that it would be in important respects desirable.[5]

This paper does not aim to address all the novel ethical issues raised by a move to learning healthcare. Instead, it aims specifically to examine the conditions under which an LHS could justifiably make participation in some forms of research *compulsory* (by way of a limited form of compulsion that I explain below). Insofar as existing discussions have addressed the question of compulsory research participation, proponents of LHSs have allowed that some research could legitimately proceed without obtaining subjects' fully voluntary consent. But they have focused their attention 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 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] Even in such cases, they have tended 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] and requiring both broad and specific forms of notification.[12] It's worth noting that, in this respect at least, these recommendations do not significantly depart from the status quo. Minimal risk research is *already* eligible for waivers of informed consent under the US regulations, and the conditions under which such waivers are available are remarkably similar to those described by proponents of LHSs.[13]

By contrast, in this paper I will argue that within the context of an LHS it could be legitimate to compel patients to participate in research presenting more than minimal risks. The specific form of "compulsion" I defend consists in making participation in research a condition on receiving care. So, I start by assuming that a provider within a hypothetical LHS has the opportunity to do some socially valuable research; I then argue that there are conditions under which they could legitimately offer participation in that research as patients' sole option for receiving care. Furthermore, I argue that this will sometimes be true even when the research in question presents more than minimal risks.

I begin by reviewing two prominent arguments in favor of compulsory research participation in the context of an LHS. I show that although the authors in question limit the research they are willing to require to minimal risk, the arguments they offer do not appear to support such a limit. I then offer my own argument in favor of compulsory research. I argue that withholding non-research options for care does not violate patients' rights, and that doing so can be both an efficient and an adequately fair means for promoting health. In determining which specific trials to make compulsory, I propose that decisions about making research compulsory are relevantly analogous to decisions about how to ration resources cost-effectively. Because there are many contending views of legitimate cost-effectiveness rationing, this leaves it somewhat open-ended what research could be made legitimately compulsory. Nonetheless, I argue that no reasonable view supports a minimal risk limit.

2. Ruth 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 patients' obligation to participate in research, Faden et al. cite both a norm of common purpose and a norm of reciprocity. 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 proportionally return the benefits we receive from others. Because medical progress is in the interest of all, and because all benefit from those who participated before, both norms support an obligation to participate.

What does this obligation require in practice? Faden et al. emphasize that proper application of their framework is not mechanical and requires the use of judgment in balancing relevant concerns. Still, they do offer some concrete direction. In particular, they indicate that the obligation to participate is properly limited by considerations of risk: patients are not obligated to participate in randomized trials of new investigational devices, though they may be obligated to participate in registries, to have their de-identified records reviewed, or to give interviews with staff. Interviews are a paradigmatically minimal risk form of research, and on the basis these examples it appears that Faden et. al. believe that the obligation to participate should be limited to research which is either minimal risk or very close to it.[15]

Yet if the obligation to participate is explained in terms of norms of common purpose and reciprocity, it is unclear why it should limited to minimally risky research. More than minimal risk research can also be a necessary means of promoting conditions which are in the interest of all. So why doesn't the norm of common purpose at least sometimes obligate patients to participate in more than minimal risk research?

Similarly, the gains we receive from past medical research are enormous. For instance, in some areas during the 18th century, smallpox may have reached an average annual mortality rate of 2,000 out of every million.[16] Because of past medical research, patients no longer face

that risk. Now compare even just that single benefit against the risks of participating in minimal risk research, where few would be willing to countenance as "minimal risk" any intervention with more than a 10 per million mortality rate.[17] Because Faden et al. do not explain how they understand proportionality, it is impossible to *demonstrate* that figures like these are out of proportion with one another. Still, I see no obvious or intuitive way in which they are.

Perhaps Faden et al. would respond that their conservative risk threshold is not derived from the obligation to participate itself, but rather arises only once the obligation to participate is balanced against others in the framework. Hence, the obligation *arises from* the norms of common purpose and reciprocity, but is still *limited by* the independent moral importance of protecting patients' interests. But this raises the question of why we should take patients' interests to so thoroughly defeat the demands of reciprocity and common purpose in this balancing process.

Perhaps in answering that question, Faden et al. could appeal to the exceptional nature of medical research as a domain requiring heightened protection.[18] But even if research exceptionalism is warranted, that, on its own, still does little to explain why compulsory research should carry a minimal or near-minimal risk limit. On the one hand, research could be exceptional, yet it could still be that patients are obligated to bear burdens much higher than those in Faden et. al.'s examples (it's just that in other non-research contexts reciprocity would demand *even more*). On the other hand, it could be that research is so exceptional that the demands of reciprocity should be voided entirely, and hence patients should not be obligated to participate in *any* research whatsoever. Without further development, Faden et. al.'s account does not explain why any particular choice here is more appropriate than any other. So, although Faden et al's account includes a minimal or near-minimal risk limit for compulsory research activities, it is unclear how their arguments support that limit.

3. Emily Largent et al.: The Veil of Ignorance

Largent et al.'s proposal for how to structure an LHS has many elements, but again we can focus on one: the "prescribed trial."[19] 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 accepting, e.g., 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 would be prescribed? Largent et al. stipulate a minimal risk limit. Anything riskier must only be "invited," a distinct category with no sanctions attached to refusal.

So Largent et al. hold that 1) some trials should be prescribed; and 2) prescribed trials should never be more than minimal risk. What justifies these two claims? They give two arguments.

The first argument is that so long as prescribed trials are limited to minimal risk, the burden of participation is negligible, yet the potential social gains are numerous and large, and this appears to be a good tradeoff. This point is sensible but limited. It justifies prescribing minimal risk trials, thereby supporting the first claim. But it does not support the second. Compare: if I say that buying a diamond ring for $20 or less is a great deal, that doesn't entail that buying a diamond ring for the slightly higher price $25 isn't *also* a great deal. Similarly, the straightforward observation that minimal risk trials seem worth prescribing does not imply that higher risk trials aren't worth prescribing as well.

By contrast, Largent et al.'s second argument is more theoretically robust and can potentially cast light on where the limits should lie: they argue that a system which prescribes trials would be selected under a contractarian approach wherein individuals are called up to choose which health system they prefer from behind a Rawlsian 'veil of ignorance' that prevents them from knowing their specific health statuses. They do not explain in any detail how such hypothetical contractors would go about making their decision, but they do suggest that they

would be willing to accept some chance of being personally subjected to minimal risk research in exchange for receiving the generally higher quality care that would be produced by an LHS.

Suppose we accept Largent et al.'s suggestion to adopt a contractarian approach, and suppose we accept their claim that those contractors would indeed choose a system that made some research participation compulsory. Would the contractors also choose to limit that compulsory research to no more than minimal risk? This depends on how, specifically, they are supposed to be making their decisions. But there are reasons to doubt that they would. Recall that Rawls himself famously thought that his veil of ignorance would lead contractors to select a form of social organization aimed at maximizing the fortunes of the worst off, regardless of the burdens that doing so imposes on other more fortunate classes. But if Largent el al.'s contractors also reasoned that way, it would seem to justify requiring people's participation in even outlandishly high-risk trials. For an extreme example, consider the possibility of running a compulsory Ebola challenge trial, wherein healthy participants are required to accept infection with the Ebola virus. If trial participants were provided with aggressive supportive care, they would still face better health prospects than many who contract the Ebola virus outside of the context of a trial.[20] So the participants being infected would not be among the medically worst off. And insofar as such a trial stood to improve the health outcomes associated with Ebola Virus Disease, it would stand to benefit some who *are* among the medically worst off. So, even shockingly high-risk research can have salutary effects according to the perspective Rawls thought hypothetical contractors would occupy.

Perhaps Largent et al. would reply that they do not agree with Rawls that idealized contractors would place absolute priority on improving the lot of the very worst off, and hence the potential consequences of doing so are irrelevant. This is reasonable enough, but they then owe us a specific account of how contractors *are* supposed to choose, if not that way. Furthermore, that account would need to simultaneously explain why such contractors *would*

choose to prescribe minimal risk trials yet *wouldn't* choose to prescribe any more than minimal risk trials. It is far from obvious what choice procedure would yield that result.

To be clear, the brief discussion above is not intended as a rigorous exploration of Rawlsian health justice, which is beyond the scope of this paper, and it is definitely not intended as advocacy for compulsory Ebola virus challenge trials.[21] Rather, the point is to illustrate that neither of Largent et. al.'s arguments for why *any* trials should be prescribed appears to justify their further claim that *only* minimal risk trials should be prescribed. Their first argument for prescribing trials is silent on what, if any, risk limit should apply. And their second argument appears to equally well justify prescribing even shockingly high-risk trials.

4. A Political Argument For Compulsory Research

We have canvassed two prominent arguments for the institution of compulsory research within the context of an LHS. Although both include a minimal or near-minimal risk limit on compulsory research activities, we found that neither appears to present a satisfying explanation for that limit. I now turn to offering my own argument for compulsory research and show that it entails rejecting a minimal risk limit.

This argument brackets questions about the "system" that will be "learning" in a hypothetical LHS. 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 just that *if* research is being done by private entities, those entities are effectively regulated such that the knowledge gained from their learning activities translates into fairly distributed future improvements in health, 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

research that, if conducted, could lead to overall health gains, under what conditions should providers be authorized to make participation in that research compulsory?

Methodologically, I approach this question by way of political philosophy, using a broad, generic framework. That framework says that the state is only permitted to adopt policies that are consistent with individual rights. And it says that from among those, it should pick those that are distributively best. That's a matter of two things: the total aggregate quantity of goods produced, and how those goods are individually distributed across people. We can refer to the first criterion as *consistency with rights,* and the next two criteria as *efficiency* and *fairness*, respectively. In sum, states should enact policies which are consistent with rights and which are efficient and fair.

In assessing compulsory research's fit with these criteria, we can avoid foundational debates. For instance, although it is a matter of great controversy which goods justice is fundamentally concerned with distributing—e.g. welfare, capabilities, or resources—all plausible views still agree that health outcomes matter in at least some instrumental or derivative respect, and hence that it is important in practice that good health outcomes be efficiently and fairly promoted.[22] Similarly, although there are diverse views on the nature and justification of norms of fairness and efficiency, it is generally uncontroversial to take greater efficiency to be at least *pro tanto* desirable, and then to take deviations from efficiency to be justified insofar as they are required by fairness. Most commonly, fairness requires deviating from efficiency in cases where doing so is necessary to adequately protecting the relatively worse-off. [23]

Under this broad framework, compulsory research can be justified by showing that there are some forms of compulsory research which are simultaneously consistent with individual rights, produce gains in efficiency, and remain adequately fair.

5. Compulsory Research Can Be Consistent With Individual Rights

Medical research often involves physical trespass on the body, and many believe that individuals enjoy a form of personal sovereignty including substantial individual rights to control what is done to and with their bodies. Whether compulsory research would necessarily violate those rights will, naturally, depend on how exactly those rights are specified as well as the form of the compulsion. Stronger forms of compulsion are more likely to be irreconcilable with individual rights. For instance, one particularly strong form of compulsion would be to perform research conscription analogously to military conscription—using a lottery to assign subjects to trials, and subsequently enforcing the selected subjects' participation using the criminal law. Although I do not find it obvious that such conscription would be morally illegitimate, the strength of compulsion involved calls for serious defense and I do not pursue it further. [24] Instead, I choose to focus only on a significantly weaker form of compulsion.

That weaker form of compulsion is a version of what Largent et al. earlier referred to as "prescribed trials." When I discuss an LHS making more than minimal risk research "compulsory," I mean making participation in that research a mandatory condition on receiving care. So, for instance, a potential patient seeking to access a therapy through an LHS might be told that if they are to receive care at all, they must consent to receive it through an ongoing or planned trial. Unlike military-style conscription, this is still consistent with allowing that the potential patient is free to decline. Some may decide to go untreated. Others may access other care sources or engage in medical tourism. Hence, the ability to opt out does preserve a fundamental measure of bodily autonomy. The body would only ever be invaded after individuals presented themselves, were adequately informed, and gave their consent.

This form of compulsion, again unlike military conscription, also appears relevantly similar to practices that are already widespread in clinical care. Providers *already* restrict the conditions under which care is offered for a variety of purposes—including purposes which may deviate from the best interests of some individual patients.[25] For instance, physicians may decide to use narrow spectrum antibiotics, despite their reduced individual effectiveness in

treating the patient before them, 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 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 about how care is offered can increase the physical discomfort patients experience, as when a trainee struggles to find a vein. They can also impose more serious risks, as when lower staffing levels lead to poorer post-surgical outcomes. Presumably, if patients were allowed to dictate the conditions of their care, many would reject being burdened in these ways: they would choose to receive procedures at their immediate convenience and from only the most experienced practitioners, to have extra staff dedicated to their stay specifically, and to be administered broad spectrum antibiotics regardless of how that might impact others. Yet they are not given these choices. Indeed, any attempt to do so would be not only inappropriate but immediately infeasible, given the potential for conflict among patients—as, for instance, if each of two patients were to simultaneously demand to immediately be taken to surgery by the same surgeon.

What these examples show is that patients' rights do not extend to controlling the conditions under which their care is offered. Rather, what patients do have is a right to authorize or decline specific interventions on their body. In making the decision whether to do so, patients have a right to the information that a reasonable person would take to be material. So, for instance, a patient has a right to know that a given facility's policy is to initiate treatment with narrow rather than broad spectrum antibiotics. And once so informed, if a patient declines to receive care under those terms, or subsequently withdraws their consent to the care at any time, then those choices must be respected. But, by contrast, what the patient does *not* have—and what it seems like they *cannot* have, given the potential for conflict among different

patients—is a right to not only reject being treated under the conditions being offered, but to demand that some others be substituted in their place.

But the form of compulsory research I described is itself just a set of conditions under which care could be offered—namely, it's a facility's choice to only offer treatment options that include research participation. It then follows from the fact that patients do not have a right to control the conditions under which care is offered that they correspondingly do not have a right against this form of compulsory research.

Importantly, this claim that patients lack an individual right to control the terms under which medical services are offered does *not* imply that the decision to offer some services and not others isn't subject to any further moral constraints whatsoever. Consistency with individual rights was only our first criterion. We still have to evaluate whether compulsory research can also be efficient and fair.

6. Compulsory Research Can Be Efficient and Adequately Fair

First, I take it as a given that any research worth making compulsory would be expected to generate aggregate health benefits, and hence would be presumptively efficient. This just as traditional research is also presumptively efficient. Of course, some traditional research is in fact ill-conceived, such that the costs and burdens associated with running it outweigh any expected benefits. But IRBs are directed to refuse approval to such research, and the same restrictions should be imposed when it comes to compulsory research. In both cases, only efficient research should be entertained.

By contrast, there is a more serious concern with how the benefits and burdens of compulsory research will be distributed across people. That is, there remains a serious concern with respect to compulsory research's fairness. Although all patients would equally retain the right to refusing to authorize research-involving treatments, under realistic conditions better-off

patients would be able to make more effective use of that right. First imagine a well-off, professional-class person with high health literacy living in a major urban center; then imagine a poor, rural, working-class person with low health literacy and limited access to transportation. Under my proposal, both hypothetical people would have the right to refuse any offer of combined research and treatment. However, one of them is much more likely to be able to pick and choose when and where they present themselves for care. So, between the two, it is the more advantaged patient that will be best able to use their right to refuse treatment to effectively protect their interests.

It is natural to worry for this reason that compulsory research would primarily enroll already-disadvantaged populations, leading to a *prima facie* unfair distribution of research burdens. This unfairness would be especially galling given that existing patterns of advantage and disadvantage in society are not themselves ethically neutral. Rather, they are frequently manifestations of injustice. Compulsory research may appear to compound this injustice, as people who have already wrongfully suffered from racism, class inequality, and other forms of oppression would also tend to be less well-positioned to avoid research enrollment, and as a result would *de facto* be required to bear yet further burdens for the benefit of others.

This is an important concern. But a full accounting of the fairness of compulsory research is significantly more complex than this initial presentation suggests. In the remainder of this section, I argue that *even if* the mechanism of compulsion I've outlined has an inherent tendency to lead to the disproportionate recruitment of disadvantaged participants, it's not obvious that this fact alone would suffice to make all possible compulsory research unacceptably unfair. Here are five reasons.

First, any reasonable analysis of the fairness of compulsory research must consider its *total* effects, and this includes assessing the way that research benefits disadvantaged populations in addition to the way in which it burdens them. Over time, the accumulation of health knowledge from compulsory research may lead to substantial improvements in health

among disadvantaged populations. From a distributive perspective, it is not clear that allowing burdens to fall on *some* disadvantaged people, in the course of producing significantly larger benefits to many *other* disadvantaged people, is overall unacceptably unfair. Furthermore, the benefits that accrue to disadvantaged populations from compulsory research could be enhanced by prioritizing research addressing conditions that disproportionately burden them. For one example, diabetes in the United States is more prevalent among non-white populations than white populations: strikingly, its incidence is roughly twice as high among Native Americans as it is among non-Hispanic whites (14.7% and 7.5%, respectively).[26] If some program of compulsory research on diabetes drew its participants disproportionately from Native American populations, then the burdens of that research would be disproportionately concentrated among Native American populations, raising the foregoing concern of unfairness. But, given the differences in incidence, the *benefits* of the research would also be disproportionately concentrated in Native American populations. This effect would be even stronger if the research resulted in a suite of interventions that worked especially well for Native American patients. Once the complete pattern of benefits and burdens is attended to, it is not obvious that the result would be less fair than the status quo. To the contrary, it is plausible that some forms of compulsory research would reduce health disparities.

Second, people who are socially and economically disadvantaged but still healthy are better off in at least some important respects than people who are gravely ill, and this suggests another way in which some forms of compulsory research could enhance fairness. Specifically, consider compelling patients with mild versions of conditions to participate in research aimed at ameliorating or preventing more serious versions of or complications from those conditions. Take HIV. HIV infection is compatible with a long, high-quality life. But it can still result in serious complications, and long-term HIV infection can cause dementia. Imagine that a study could be done on healthy people with well-controlled HIV infections, with the goal of contributing to the development of better techniques to prevent HIV-associated dementia. Suppose such a study

were made compulsory, and then went on to enroll mostly racially and economically disadvantaged participants. That would result in a regrettable distribution of research burdens. Nonetheless, a healthy person with a well-controlled HIV infection, even if racially and economically disadvantaged, is still likely to have more relevant capabilities and to enjoy a higher quality of life than a person suffering from an advanced case of HIV-associated dementia. If compulsory research were to impose burdens on some socially and economically disadvantaged but relatively healthy people, while also contributing to an eventual reduction in HIV-associated dementia, then the overall social distribution of relevant life goods would be made more rather than less fair. This is another way in which compulsory research could improve the fairness of the status quo even under the supposition that the relatively healthy people it burdened would tend to be disproportionately drawn from already socially and economically disadvantaged classes.

Third, an LHS might make active efforts to promote fair burden sharing. For instance, it might reserve its most burdensome research for its best-performing care centers. Would affluent patients then abandon those centers rather than submitting to participation? Not necessarily, if they were still to perceive the center's baseline care to be sufficiently desirable. If so, it would give the LHS another way of strategically burdening relatively advantaged populations in the course of performing research that, given appropriately selected research questions, could potentially benefit all.

Fourth, an LHS might attempt to offset research burdens with inclusion benefits, most obviously by waiving fees and/or paying patients who participate in research, but also by providing perks like preferential scheduling for appointments or preferential assignment to private rooms. By reducing the net burden of participation, such measures could also reduce concerns of unfairness in the distribution of research burdens.

Fifth and finally, at least some unfairness can be justifiable. Allow, just for the sake of argument, that it is impossible to design a compulsory study that was not at least somewhat

unfair in its distribution of research burdens. Even if that were so, it may be that said unfairness can nonetheless be justified by the expected health gains from the research. This is just a consequence of the generic political approach we outlined earlier, which treats both efficiency and fairness as relevant considerations. The person who seeks to reject compulsory research on fairness grounds must not only hold that it is inevitably unfair, but that it is inevitably so extremely unfair as to be unacceptable however large the benefits in health it might bring. Interpreted as an exceptionless claim about all possible studies that might be made compulsory, I find this implausible, particularly in light of the foregoing arguments.

For these reasons, I conclude that in addition to violating no rights and being presumptively efficient, compulsory research can in at least some cases be adequately fair.

7. Compulsory Research As Cost-Effective Healthcare

Compulsory research *can* be efficient and fair. But that is not to say that always is. The previous section described specific ways of designing compulsory research to enhance fairness, but these may not always be available, effective, or appropriate. Even if they are available, effective, and appropriate, they may not adequately counter the unfairness-producing tendency of compulsory research to disproportionately recruit the disadvantaged. Indeed, even if the previous section succeeded in arguing that there are conceivable circumstances wherein compulsory research would be adequately fair, a skeptical reader may be left with the concern that, in absence of a clearer method for practically identifying those circumstances, the difficulty in responsibly deciding when to license compulsory research would be too great. Fortunately, such worries can be addressed. There is a clearer method for identifying when exactly a study should be made compulsory, and, furthermore, this method can explicitly incorporate requirements of adequate fairness. That method is to treat the decision to impose compulsory research as a special case of cost-effectiveness rationing.

To illustrate this proposal, start with a standard and relatively uncontroversial example of cost-effectiveness rationing: the ongoing decision of a health system to not build new, top line facilities in sparsely populated rural areas, on the grounds that anticipated low utilization would render such facilities cost-ineffective. As a result, fewer health services are available in those areas. The justification is that the savings gained can be used elsewhere to produce yet larger health benefits, so rationing efficiently promotes overall health. Yet the choice raises worries about fairness. Maintaining a lower density, variety, and/or quality of services in rural areas does not burden everyone equally. It burdens rural populations. It can also entail seriously elevated risks, as in areas where long travel times mean emergency care is often too late. And in many contexts, rural populations will also be relatively disadvantaged, hence the already burdened are further burdened.

As before, this is especially galling when the patterns of relative advantage being reinforced by a resource allocation scheme are not themselves ethically neutral. For instance, imagine that the rural population in question is a Native American population living on tribal land. In that case, the low density and remoteness of that patient population may be a result of state-sponsored forced relocation. Using those features of the situation to then to justify a sparser allocation of health services may appear to objectionably compound this underlying injustice.

Yet despite this worry about the potential for rationing to compound unfairness, there is no reasonable alternative. Health resources are limited, and it would be absurd to allocate them in complete disregard of their effects: a health system *shouldn't* overinvest in a facility located in a remote rural area where it's destined to only ever serve low volumes of patients. Instead, the relevant question is how to temper the efficient production of health in light of considerations of fairness. So although siting a flagship facility in a rural outpost wouldn't make sense, a health system may still be reasonable in paying higher rates for rural services than would be dictated by a strictly efficient resource allocation. It may also be willing to pay higher rates when

evaluating services to populations which have a history of being targeted by violence and oppression. Exactly how much more it should be willing to pay, and, more generally, how it ought to carry out this balancing of fairness against efficiency in resource allocation is the subject of theories of justice in cost-effectiveness rationing. This is a subject on which there are rage of reasonable competing views.

The point of sketching this example is not to argue in favor of any particular funding levels for rural health services, nor to attempt to resolve the correct response to the ways in which patterns of current social advantage reflect both present and historical injustice. Rather, the point is to demonstrate the striking parallel between cost-effectiveness rationing and compulsory research. Both involve the same mechanism: they involve deciding to offer what may be less desirable, more burdensome, and even riskier treatment options. They involve structuring offerings in ways that leave some patients facing prospects which are overall worse than they theoretically could be. They do so for the same reason: to more efficiently promote health. And they raise the same worry: that their efficiency must be balanced against their potential unfairness. These similarities are sufficiently robust that compulsory research can itself be evaluated as a form of cost-effectiveness rationing. Indeed, it *should* be so evaluated, and, whatever the correct theory of justice in cost-effectiveness rationing is, it should be applied the same way to each. When the same underlying issues present themselves across multiple contexts, they should be resolved consistently in each.

For a very simple example of how cost-effectiveness evaluation could guide decisions about compulsory research, start by supposing that there is some drug which is already in clinical use, but where there remains an important question about whether it crosses the blood-brain barrier. Hence, there is interest in running a pharmacokinetic trial where, in addition to receiving the drug as normal, participants would be subjected to some number of lumbar punctures (LPs) for purely research purposes. Could this trial reasonably be made compulsory? I propose a two-step evaluation.

In the first step, imagine that the trial, including all its elements, were a single treatment, and evaluate that fictional single treatment along cost-effectiveness lines. In this case, that involves imagining that the LPs were not research interventions, but instead a necessary part of a therapeutic package. LPs are mostly safe but uncomfortable and pose a small risk of severe, lasting headache. The badness of this discomfort and headache risk can be quantified, just as discomfort and headache risk would be quantified in the cost-effectiveness assessment of any treatment that caused them. At the same time, the trial would generate valuable knowledge. That value can also be quantified, this time in terms of how much the health system thinks is reasonable to pay for it. The cost effectiveness of the trial, taken all together, would then be that of the base drug, but supplemented on the one hand with the burdens of the LPs added on *as-if* unavoidable noxious side effects, and on the other with the value of the knowledge gained added on *as-if* a special discount reducing the cost of provision.

In the second step, compare the resulting assessment against the cost effectiveness of providing standard clinical care. If running the trial is less cost effective than providing standard clinical care, then it should not be prescribed. But if it is superior, then the choice the health system faces is relevantly like the one it has whenever it must decide whether to offer a treatment option that is less desirable for many individual patients yet also more cost-effective. So it should deploy whichever reasonable theory of justice in cost-effectiveness rationing it favors. If that theory would recommend rationing people to the trial when the trial is imagined as if it were a single treatment, taking into account both the extra burdens it imposes on patients as well as the value it generates, then making participation in the trial compulsory is ethically appropriate.

There are three things to note about this proposal.

First, this method refers the question of whether to prescribe any specific trial to adjudication by a background theory of justice in cost-effectiveness rationing. Not only does this give a systematic answer to when studies should or should not be prescribed, it does so in a

way that naturally addresses worries about fairness, given that theories of justice in cost-effectiveness rationing already routinely take into account fairness by prioritizing interventions to the less well-off. And, as noted, attention to injustice and oppression can be incorporated via similar means, by similarly prioritizing the correction of health problems that derive from or compound unjust underlying social arrangements. In the same way that a health system might "over-weight" the benefits of correcting rural health problems relative to what it would otherwise pay, in order to ensure that rural populations are not unfairly under-served, so too could an LHS "over-weight" the badness of research-related burdens when those burdens would be foreseeably concentrated within a disadvantaged population or when their occurrence within a population would reflect or compound existing injustices. This would effectively make the LHS less willing to go ahead with compulsory research, even valuable research, in proportion to the degree to which the research-related burdens would be unfairly distributed. Given that such fairness-related correction factors must already be decided on in order to make routine rationing decisions about the provision of clinical care, those same correction factors can also be used to guide decisions to compel research participation.

This approach is consistent with placing an extremely high priority on fairness, given that the correction factor used to overweight the "badness" of unfairly-distributed research-related burdens could be set arbitrarily high. That is not to say that it would be a good idea to do so, as higher correction factors will also force larger sacrifices in efficiency. Rather, some balance must be struck, and, again, where precisely that balance should be struck is not a question this paper aims to answer. Instead, the point here is that concerns with fairness do not support a blanket rejection of compulsory research, nor do they undermine this paper's identification of the ethics of compulsory research with the ethics of cost-effectiveness rationing. Rather than such concerns leading readers to reject the central theses of this paper, instead they should lead such readers to accept the framework of the paper and then go on to argue within it that the

appropriate fairness-orweighting factors should be set at the higher end of the reasonable range.[27]

The second thing to note about my proposal is that the example given to illustrate it was very simple. Testing of a novel drug, by contrast, will introduce uncertainty that does not exist with regards to a standard therapy and a few LPs. I believe that it is possible to continue relying on the guiding thought that compulsory research should be treated as-if it were a fictional therapy being evaluated for cost effectiveness while extending the reasoning to cover these and other complexities. In this case, that guiding thought would call for attention to situations where health systems must decide between an array of offerings with different associated costs, but whose risks and benefits, perhaps because not yet adequately studied, are still substantially uncertain. But carrying out these extensions is beyond the scope of this paper.

The third thing to note about my proposal is that allowing that there are a range of reasonable views on appropriate cost-effectiveness rationing leaves the practical question of what research should be made compulsory somewhat open-ended. Still, the identification of the ethics of compulsory research with the ethics of cost-effectiveness rationing is already itself a form of substantial theoretical progress: it demonstrates that there is no need for a radically new ethical framework to govern learning healthcare. What's more, although the concrete implications of the present proposal are admittedly open-ended, there are still at least some substantive conclusions we can draw regarding compulsory research. Returning to the question we started with, we are now in a position to reject a minimal risk limit. To illustrate: no one would (or should) think that a safer formulation of a drug that leads to 11 per million fewer fatal allergic reactions is worth paying *literally any price* for. All legitimate views of cost-effectiveness rationing—all that lie within the "reasonable range" I've been referring to—will agree that at some point, considerations of cost will override that magnitude of a reduction in risk. But as we might recall, few allow as "minimal risk" any research intervention which would lead to a greater than 10 per million increase in mortality. Given that more than minimal incremental increases in

risk can be justified by the logic of cost-effectiveness rationing, provided the savings are significant enough, it follows that there will also be more than minimal risk potential trials which would be appropriate to make compulsory, provided the relevant knowledge is valuable enough.[28] If the parallel to cost effectiveness rationing is accepted, it follows that a minimal risk limit cannot be defended.

8. Implementing Compulsory Research

Before concluding, consider an objection. Traditional cost effectiveness analysis balances the contemporaneous health benefits of an intervention against the cost of providing it. But the costs and benefits of compulsory research do not occur at the same time, or even necessarily in 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 greater health benefits to other patients *now.*

I doubt this is true. Notice that the aforementioned cases of restricting antibiotic access and physician training also impose burdens on patients in the here and now in exchange for future benefits which are also not guaranteed to manifest in the same system. Yet they seem clearly defensible. Still, I am willing to allow this objection for the sake of argument.[29]

But then consider the following arrangement. Imagine a health funding agency that starts by evaluating a large portfolio of research proposals: it estimates, for each, the health gains associated with performing the research and how much those gains are “worth” under its institutional view of health promotion. It then uses those estimates to determine the amount it would be willing to spend to see each trial conducted, and it publicizes those amounts. A hospital, insurance network, or other provider would then have the option to choose to run one of those trials, and, in return, to receive funding in the amount reported. In deciding whether to do so, they would deploy their favored view of cost-effectiveness rationing, taking into account

both the burdens to patients yet also the benefits of the additional funds. If they decided to prescribe participation, patients at that system would then have the choice of whether to access service through the trial or instead attempt to find healthcare through other channels.

For all three parties to this situation, their decisions under the arrangement I just described look much 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 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 the way they already do--with the difference being that their actions are brought into an alignment that systematically produces health gains from research. When we can produce such positive alignments, we should.

9. Conclusion

Existing calls for the implementation of a learning healthcare system have been cautious with respect to compelling research participation and have tended to stipulate a minimal or near-minimal risk limit for compulsory learning activities. I have argued to the contrary. So long as the relevant form of compulsion consists solely in withholding non-research options for receiving care, that withholding does not violate any individual's rights. Furthermore, doing so can be both efficient and adequately fair. Because the decision to make research compulsory is relevantly similar to the decision to ration care, theories of justice in cost-effectiveness rationing can be

used to further specify when to use compulsion. When applied, these theories will not support a minimal risk limit.

* * *

References:

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.
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.
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?"
14. Faden et al., "An Ethics Framework"
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. Edwardes, Edward. *A Concise History of Small-pox and Vaccination in Europe.* (London: H.K. Lewis, 1902), 14, https://curiosity.lib.harvard.edu/contagion/catalog/36-990060994170203941
17. 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. See also 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. For sustained critique of research exceptionalism, see A. Wertheimer, *Rethinking the Ethics of Clinical Research: Widening the Lens* (Oxford: Oxford University Press, 2011). For a defense, see J. Wilson, D. Hunter, "Research Exceptionalism," *American Journal of Bioethics* 10, no. 8 (2010):45-54.
19. 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?"

20. 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.

21. It is worth noting that the most prominent Rawlsian account of health justice does not locate the distribution of health under the difference principle, but instead under fair equality of opportunity. Re-running this section's discussion in terms of fair equality of opportunity strikes me as more complex, but no more promising when it comes to recovering a minimal risk limit. N. Daniels, *Just Health: Meeting Health Needs Fairly* (New York: Cambridge University Press, 2007).

22. Within a distributive paradigm, 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 health, that itself discredits them. Similarly, it is not obvious how relational egalitarians should treat health. Nonetheless, they should hope to have some explanation for the significance of health, if their view is to be plausible.

23. Although prioritarians and egalitarians disagree over what fundamentally grounds the desirability of a distribution, as an extensional matter both will be willing to trade efficiency for improvements to the lot of the worst off. And even utilitarians, who do not share their distinctly distributive concerns, will typically agree in light of the diminishing marginal utility of many resources.

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

25. 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 still emphasize the 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.

26. "Statistics About Diabetes," American Diabetes Association, accessed March 6th, 2017. https://www.diabetes.org/resources/statistics/statistics-about-diabetes.

27. This is one important way in which fairness can be built into cost-effectiveness assessments. In highlighting it, I do not mean to imply that there could not be yet further, additional requirements. For instance, one reviewer wondered whether compulsory research should be governed by a responsiveness requirement analogous to traditional clinical research. I am open to that possibility. In keeping with the guiding analogy here, the present approach would advise resolving this question by first asking whether it is a condition on legitimate cost-effectiveness rationing that communities which are subject to reduced offerings must directly benefit from the savings produced by that reduction, then taking that answer, whatever it is, to presumptively apply to compulsory research as well.

28. An anonymous reviewer requested a concrete example of a more-than-minimal-risk study which could legitimately be made compulsory, and, ideally, which could not be carried out in any other way. Yet it is complicated to give such an example while remaining within the bounds of the paper. First, which studies are minimal risk is hard to pin down in practice (different IRBs decide differently). Second, it is, as noted in the main, also controversial when cost-effectiveness rationing is legitimate, and hence controversial when compulsory research can be justified by its similarity to cost-effectiveness rationing. So the present paper argues that everyone who has a reasonable view of cost-effectiveness rationing should be willing to agree that there is some study which is more than minimal risk and also can be legitimately made

compulsory, but allows that different people with different background views may still disagree amongst themselves on what such a study looks like.

29. An anonymous reviewer raised a similar version of this same concern about the justification of cost-effectiveness rationing, but they put it in terms of institutional goals: they argued that the mission of an academic hospital is to support the current health of the community, not to try to improve its future. I believe the example of physician training shows that this present-time-centered view of the mission of care is unsustainable, but, even if it were to be granted, the scheme I outline in the rest of this section would still enable the present burdens of research to be offset with present benefits.