

Wenner, D.M. (2017). "The Social Value of Knowledge and the Responsiveness Requirement for International Research," *Bioethics* 31(2): 97-104 (doi:10.1111/bioe.12316)

This is a pre-print of a published manuscript. When citing, please refer to the final, published version at <http://onlinelibrary.wiley.com/doi/10.1111/bioe.12316/full>.

Address correspondence to:
Danielle M. Wenner
Department of Philosophy
Carnegie Mellon University
155C Baker Hall
Pittsburgh, PA 15219
dwenner@andrew.cmu.edu

The Social Value of Knowledge and the Responsiveness Requirement for International Research

Introduction

Ethicists have long recognized that two necessary features of ethical research are scientific validity and social value.¹ This principle is enshrined in not only the publications of ethics scholars, but also several international guidance documents related to the ethical conduct of research.² Yet despite a significant literature surrounding the validity component of this dictate, until recently little attention has been paid to unpacking what the social value component might require.

My purposes in this paper are twofold: First, to introduce a framework for assessing the social value of research, and in particular, for determining prospectively whether a given research program is likely to have social value of the kind necessary to fulfill the social value requirement. And second, to illustrate how this framework can provide a clearer account of the responsiveness requirement, an ethical dictum that is oft-repeated but whose content and value as a guideline in international research remain contested. Although this conception of social value is introduced in the context of the responsiveness requirement, the account offered is capable of standing alone as a tool for the assessment of the social value of research, and intended to do so. As such, it can be usefully applied to discussions about the value of comparative effectiveness trials, the continued proliferation of me-too drugs and the research done to develop them, the demand for public access to study-generated data, and persistent deficits in the publication of negative results, among others.

¹ E.J. Emanuel, D. Wendler, & C. Grady. What Makes Clinical Research Ethical? *JAMA* 2000; 283:2701-2711.

² See, for example, Council for International Organizations of Medical Sciences (CIOMS), "International Ethical Guidelines for Biomedical Research Involving Human Subjects," (2002).

In Section 1 I introduce the responsiveness requirement, survey the primary justifications offered for this requirement in the literature, and explore the shortcomings of existing accounts of the requirement. Section 2 introduces an account of the social value of knowledge (SVK) that is borrowed from decision theory, and suggests that a helpful way of characterizing responsiveness is as a demand that clinical research be capable of generating knowledge that is valuable to the host community. Finally, Section 3 outlines what is necessary to fully specify SVK and introduces some of the hurdles to full specification, but also indicates how the framework can function as a threshold condition to inform responsiveness even when under-specified.

Section 1: The Responsiveness Requirement

Clinical research conducted in low- and middle-income countries (LMICs) plays a growing role in the research and development of new biomedical interventions. Given disparities in wealth and access to healthcare between higher and lower income settings, many worry that LMIC populations are vulnerable to exploitation in clinical research.³ Moreover, some worry about the “10/90 gap”, that less than 10% of global health research resources are utilized in addressing the health needs of LMICs, where greater than 90% of preventable deaths occur.⁴ Proponents of the responsiveness requirement suggest that by limiting the research that can be conducted within LMIC populations to that which is “responsive” to local health needs, host communities can be protected against exploitation by research sponsors from high-income countries. Moreover, they suggest that such limits may assuage concerns about the 10/90 gap.

³C. Grady. Ethics of International Research: What Does Responsiveness Mean? *Virtual Mentor* 2006; 8: 235-240.

⁴The phrase “10/90 gap” refers to the disparity in 1990. Since then the gap has closed somewhat, but LMICs still account for greater than 80% of preventable deaths worldwide.

When critics complain about exploitation in clinical research, they are typically not claiming that trial agreements were not entered into voluntarily.⁵ Even when all parties consent to research, critics can claim that host communities are exploited. Moreover, the charge of exploitation need not imply that host communities aren't benefiting from research; often they are. Rather, the charge of exploitation is a complaint about the resulting distribution of the social surplus generated by the interaction: it is somehow unjust or unfair.⁶ Typically, the claim is that the exploited party receives a benefit that is disproportionately small as compared to the benefit to the exploiting party⁷ or the needs of the exploited party.⁸

Insofar as proponents of a responsiveness requirement are seeking to reduce exploitation, then, the claim must be that responsiveness serves to ensure that research transactions result in greater benefits for host communities, or that some research transactions that don't provide sufficient benefits to host communities are thereby prevented. And since responsiveness specifically constrains what kinds of research can be conducted within LMICs, the claim must be that the responsiveness requirement will serve these goals in some way related to the products of research. Similarly, if proponents of responsiveness seek to address concerns about the 10/90 gap, this must be a claim about the kind of research that is being conducted in LMICs. Because responsiveness functions as a *limit* on the research that can be conducted, the claim that the requirement will incentivize research that addresses health deficits would be difficult to substantiate, but is not conceptually incoherent: while some sponsors or researchers may be

⁵ There is nothing conceptually incoherent about exploitation that is not morally wrong: I might exploit my tennis opponent's backhand without thereby wronging her. The focus of the charge of exploitation for my purposes here is morally wrongful exploitation. For ease of exposition, I shall simply refer to this as "exploitation" and the reader may assume that I am referring to that subset of exploitative interactions that are morally problematic.

⁶ D.M. Wenner. Against Permitted Exploitation in Developing World Research Agreements. *Dev World Bioeth* 2016; 16:36-44.

⁷ C.f. A. Wertheimer. 1996. *Exploitation*. Princeton: Princeton University Press.

⁸ C.f. R. Sample. 2003. *Exploitation: What it is and Why It's Wrong*. New York: Rowman & Littlefield.

discouraged from conducting research in LMICs at all, some may instead conduct more responsive trials. A more conservative claim would be that by limiting research in LMICs to that which is responsive, reinforcement of the 10/90 gap can be precluded by ensuring that any research that *is* conducted in LMICs will erode background health deficits rather than exacerbate them.

Although research participants are individually vulnerable to exploitation, it is unclear how a constraint such as responsiveness could serve to reduce exploitation of individuals as such, at least directly. Responsiveness relates to the content of the research question and how resulting benefits should be relevant, and in general research is not conducted with the intention of benefiting participants via the investigational product. Research is conducted to resolve uncertainty about the likely effects of various interventions or their components. Responsiveness thus seems better suited to address concerns about exploitation of host communities, who can more naturally be construed as the beneficiaries of research outputs. That said, we might consider responsiveness as indirectly addressing the exploitation of research subjects, insofar as it functions to ensure that after a trial, participants have access to a health system that has been improved by the research in some way. In what follows, I take for granted that responsiveness, if able to reduce exploitation, does so by reducing exploitation of host communities rather than individual participants. This leaves open the possibility that additional conditions are necessary to protect research subjects themselves from exploitation in developing world research.

Objections to existing conceptions of responsiveness are well-rehearsed in the literature. But the orientation towards reducing exploitation and addressing the 10/90 gap, once fully enunciated, provides further reason for skepticism about those accounts. Both considerations suggest that in order to do the work that responsiveness is intended to do, the requirement cannot

function merely to limit research to locally prevalent conditions. After all, research on locally prevalent conditions might nevertheless fail to benefit host communities via the products of research. Consider the well-known Surfaxin trial, in which Discovery Labs sought to test a new treatment for infant respiratory distress syndrome (RDS) against placebo in several Latin American settings. Given its high cost, Surfaxin was unlikely to be marketed to the test populations, who absent the trials lacked access even to respirators with which to treat RDS in infants.⁹ This trial might be construed as “responsive” to local health needs, insofar as it addressed a condition of local concern. But given the lack of intention to market Surfaxin locally or otherwise ensure its availability to host populations, the trial’s “responsiveness” on this conception does nothing to cut against exploitation or the 10/90 gap.

Given this concern, many guidance documents additionally require that study interventions proven efficacious be made reasonably available within host populations after a study’s conclusion.¹⁰ Reasonable availability seeks to address concerns that research will be conducted on locally prevalent conditions when resulting interventions will be cost-prohibitive for local use. This seems to capture the intuition behind responsiveness: that research conducted in LMICs should contribute to improving local access to needed healthcare. But demands for reasonable availability assume that researchers and sponsors can ensure local approval for dissemination of new interventions after a trial, when in reality local regulatory structures can be difficult and time-consuming to navigate. A more plausible interpretation of reasonable availability might require that sponsors and investigators attempt to negotiate these hurdles in

⁹ J.S. Hawkins & E.J. Emanuel. 2008. Case Studies: The Havrix Trial and the Surfaxin Trial. In *Exploitation and Developing Countries: The Ethics of Clinical Research*, J.S. Hawkins & E.J. Emanuel, eds. Princeton: Princeton University Press: 55-62.

¹⁰ See, for example, Council for International Organizations of Medical Sciences (CIOMS), *op. cit.* note 2; National Bioethics Advisory Commission, “Ethical and Policy Issues in Research: Clinical Trials in Developing Countries,” (2001).

advance, but even if approval is obtained, it is unclear who should have access to the interventions, via what mechanisms, and for how long.

Reasonable availability also provides no guidance for studies that don't lead directly to new interventions, such as observational or epidemiological studies or early phase research whose causal contributions to the development of new treatment and prevention modalities is murky, at best.¹¹ Given the necessary redundancy in the production of biomedical evidence, a functional conception of responsiveness should be able to discriminate between responsive and non-responsive research even when a particular clinical trial cannot be credited with leading directly to a positive assessment of an intervention's effectiveness and there is therefore no product to be made available at a trial's conclusion.¹²

A more useful interpretation might demand that any *knowledge gained* be made reasonably available, and in fact this interpretation is consonant with the language in the 2001 version of the CIOMS Guidelines, which states that "Any intervention or product developed, *or knowledge generated*, will be made reasonably available for the benefit of that population or community."¹³ But consider the Surfaxin trial again. Even had host communities been granted access to the information gained in the trial, it would have generated little value for them given competing claims on already inadequate local health resources. This interpretation would therefore do little to assuage concerns about exploitation or underlying health deficits. Ultimately, even with the additional requirement of reasonable availability, responsiveness to health needs is unable to do the work it is intended to do.

¹¹ D.M. Wenner. The Social Value of Knowledge and International Clinical Research. *Dev World Bioeth* 2015; 15: 76-84.

¹² Participants in the 2001 Conference on Ethical Aspects of Research in Developing Countries. *Science* 2002; 298: 2133-2134.

¹³ Council for International Organizations of Medical Sciences (CIOMS), *op. cit.* note 2, emphasis added.

A stronger conception of responsiveness requires research to be responsive to health *priorities* of host communities, seeking to ensure that studies have the potential to address important health deficits that are not able to be met given existing local resources.¹⁴ Application of this interpretation seems relatively straightforward where local health priorities have been enumerated, but where they are unspecified or too vague to clearly signal which research addresses them, we are left without a benchmark for assessment.¹⁵ In such contexts, it is an open question how best to determine what constitutes a local priority, a question complicated by potential difficulties in obtaining complete and accurate epidemiological data. Absent specification of research priorities by local authorities, then, some clear method of identifying appropriate research which is consonant with the motivations grounding responsiveness is required.

One approach might specify priorities relative to the incremental health benefits of new interventions. Objective measures of disease burden such as Quality-Adjusted or Disability-Adjusted Life Years combined with prevalence measures can highlight the greatest sources of mortality and morbidity within a population when epidemiological data is available, and can be used to quantify the “health gain” or differential effectiveness of new interventions as compared to established alternatives or no treatment. Metrics reflecting expected uptake of new interventions can be used to predict their overall health impact, assigning value on the basis of the expected direct, causal impact of a study on the development of a new, effective intervention.

¹⁴ A.J. London. 2008. Responsiveness to Host Community Health Needs. In *The Oxford Textbook of Clinical Research Ethics*, E.J. Emanuel, et al, eds. Oxford: Oxford University Press: 737-744. The draft 2016 revision of the CIOMS guidelines, available <http://www.cioms.ch/>, favors this interpretation.

¹⁵ Shah et al. report that Kenya’s stated research priorities include “research on non-communicable diseases as well as communicable diseases.” From this we might infer that Kenya desires a diverse research portfolio, but little else regarding which diseases or conditions would be appropriate for study within this population. S. Shah, R. Wolitz, & E.J. Emanuel. Refocusing the Responsiveness Requirement. *Bioethics* 2013; 27: 151-159: 156.

This approach might be able to accommodate some early-phase research by assessing the likely health impact of expected later-phase outputs, and then discounting that value based on the expected time to market.¹⁶ But the development of new biomedical products is not linear in a way that allows the health impacts of all research outputs to be straightforwardly estimated. Many research outputs have unclear causal links to the future development of specific new interventions, but nevertheless can perform valuable roles such as informing new or revised research programs or providing new insights into biologic functions or novel delivery mechanisms.¹⁷

Moreover, health impact measures assign minimal value to research addressing rare diseases or highly prevalent conditions with lower disease burdens. We might think that research within LMICs ought to focus on having the greatest potential health impact, but it is similarly important to respect the autonomy of host communities. Ideally, the principle informing responsiveness will leave room for communities to prioritize their health research interests in some ways that stray from strict impact measures in accordance with local social will. Otherwise, we risk depriving communities of their input in balancing the competing social values at stake in local priority-setting.

Finally, absent accurate epidemiological data, objective impact measures are difficult to apply. What is needed is a conception of responsiveness that can offer clear guidance without accurate epidemiological data, that can accommodate research not leading directly to new interventions, and that can allow for autonomous local priority-setting while still ensuring that

¹⁶ D.J. Casarett, J.H.T. Karlawaish, & J. Moreno. A Taxonomy of Value in Clinical Research. *IRB* 2002; 24: 1-6.

¹⁷ J. Kimmelman. 2010. *Gene Transfer and the Ethics of First-in-Human Research: Lost in Translation*. Cambridge: Cambridge University Press.

trials conducted in LMICs are capable of generating benefits to host communities via the outputs of research.

Section 2: The Social Value of Knowledge

The rest of this paper introduces a framework for the assessment of the social value of research, and goes on to demonstrate how this framework can be used to inform the responsiveness requirement. Above, I argued that when critics claim international research is exploitative, they are frequently making a claim about the distribution of the benefits from clinical research. Research often brings much-needed benefits to low-income populations such as access to basic healthcare, infrastructure development, or training of medical personnel. But given the urgency of need that can exist, when communities are effectively competing with one another to attract research they can quickly race to the bottom with respect to the benefits they'll accept in exchange for hosting trials.¹⁸ The charge of exploitation suggests that host communities end up receiving less than their fair share of the social surplus of these interactions.

One way to ensure host communities receive more of that surplus is to require benefits that are not one-off but continuing. The claim of responsiveness proponents appears to be that by ensuring host communities receive a benefit of a certain *kind*, namely, benefits from research outputs which will have an ongoing impact on local health systems, we might ensure host communities receive *greater* benefits. Moreover, those benefits will be of a kind that reduces health disparities. What is needed is a less problematic way to characterize what kinds of research outputs can play this role.

¹⁸ A.J. London & K.J.S. Zollman. Research at the Auction Block: Problems for the Fair Benefits Approach to International Research. *Hastings Cent Rep* 2010; 40: 34-45.

The conception of responsiveness that I propose is a special instance of the social value requirement. In particular, I interpret responsiveness as a requirement that research conducted in LMICs be aimed at producing knowledge which is socially valuable specifically to its host communities. Because of the specialized nature of research-generated knowledge, in order for that knowledge to be of *social* value, it must be of instrumental use in some way that generates benefits for society.¹⁹ Responsiveness is thus a function of the relation between the question being asked in a particular trial and the usefulness of answers to that question to local populations.

Suppose Alice is a policy-maker facing some policy decision, and suppose further that she faces some discrete number of policy choices along with a set of outcomes or consequences that might result from whatever policy she chooses. She might be considering local school vaccination policies, for example, and considering likely outcomes with respect to what proportion of students get vaccinated, potential health outcomes, and how prohibitive policies might be for individual families, whether due to personal beliefs, economic barriers, or other considerations. Since it is possible for two different policy choices to generate the same outcome, assume also that society values a given outcome the same regardless of the policy choice that leads to it.

We can use utility functions to represent how society values different outcomes: Assume society places a value of 100 on outcome o_1 and 30 on outcome o_2 . In our example, these values might represent some proportion of students vaccinated, some associated level of morbidity, and some quantity of barrier to student enrollment. Given the information Alice has, she will choose

¹⁹ M.G.J.L. Habets, J.J.M. van Delden, & A.L. Bredenoord. The Social Value of Clinical Research. *BMC Med Ethics* 2014; 15:66.

the policy with the greatest expected utility, which she figures according to the value society assigns each outcome and the likelihoods of each option leading to particular outcomes. If Alice thinks policy A has a 25% chance of generating outcome o_1 and a 75% chance of generating outcome o_2 , and policy B has a 35% chance of generating outcome o_1 and a 65% chance of generating outcome o_2 , then she'll assign expected utilities of $(.25 \times 100 + .75 \times 30)$, or 47.5 to policy A and $(.3 \times 100 + .7 \times 30)$, or 51 to policy B. With her current information, Alice should choose policy B to maximize expected social utility.

But new information can cause Alice to revise her assessment of how likely the possible outcomes are under different policy choices, causing a revision of her expected utilities. For example, she might now think that in fact the likelihood of outcome o_1 is 35% under policy A, and outcome o_2 65%. (Maybe she's seen new data about parental pushback to policies prohibiting enrollment without vaccination.) This would change policy A's expected utility to 54.5 $(.35 \times 100 + .65 \times 30)$. In the absence of any information revising likelihoods attached to policy B, Alice should change her decision to A since A now has the highest expected utility of available options.

We can generalize this as follows: whenever new information causes a substantial enough revision of the probabilities assigned to different outcomes under some policy option(s) to change which policy has the highest expected utility, a policy maker should alter her choice. In such cases, the difference between the *ex ante* maximum expected utility (the expected utility of the most attractive decision absent the new information) and the *post hoc* highest expected utility (the expected utility of the most attractive choice with the new information) represents the value

of the knowledge gained.²⁰ If new information doesn't change the preferred option, it is not instrumentally valuable on this account.²¹

The corresponding conception of responsiveness requires that the knowledge likely to be gained in a trial have the potential to inform local decision-making in this way. The relevant decisions that might be informed by trial-generated knowledge include those regarding local health policy or resource allocation, local research priorities, and the funding decisions of NGOs working locally. Hereafter, I refer to this as the “SVK Framework”. In the next section, I explain what we need to know to use this framework and how it can offer clear guidance in cases where alternative conceptions of responsiveness cannot. In answering these challenges, this framework is better suited to play the role that responsiveness proponents envisage for the requirement. I also indicate where there are difficulties in fully specifying social value, and how the framework can be used to inform responsiveness as a heuristic in the absence of full information.

Section 3: Required Inputs & Remaining Challenges

Most basically, the prospective assessment of clinical research under this framework requires three inputs: what might be learned from a given trial, how probable it is that we might learn it, and how that information might alter the decision-making of relevant stakeholders. This last question requires identification of the community to whom research should be valuable. Those who accept the social value requirement but reject responsiveness will refer to the global community here, and I should point out that the SVK Framework can function at the global level to inform judgments of social value *simpliciter*. That said, this point ultimately makes contact

²⁰ I.J. Good. On the Principle of Total Evidence. *Br J Philos Sci* 1967; 17:319-321.

²¹ The implication is not that such knowledge is not valuable *per se*, but that it is not socially valuable in the manner required to meet the social value or responsiveness requirements.

with the motivations behind responsiveness. If trials are only limited to questions of value to the global scientific community, this leaves unaddressed concerns that LMIC populations will bear a disproportionate share of the burdens of biomedical research while reaping few of the long-term gains in improved health.

For any account of responsiveness, then, the social context to which responsiveness should be indexed will be a central question. One suggestion is to link responsiveness to national decision-making.²² This makes sense insofar as national officials may be best informed about health needs and available resources within a country. However, this approach seems to assume that national health priorities are consonant with local ones. If the concern motivating responsiveness is that the health and well-being of high-income populations ought not to be disproportionately benefited by research conducted on low-income populations who don't themselves receive relevant health gains, the same worry speaks to the distribution of health resources within an LMIC. If central authorities are more in touch with urban population centers but want to authorize research in poorer, rural areas to which medical advances may not find their way, this is problematic on precisely the same grounds.

A better way to conceptualize the relevant community references health systems. In contexts similar to the UK where a National Health Service distributes access to healthcare in a largely egalitarian fashion, there may be little difference between local and national health systems. But where health systems are fragmented, responsiveness should consider the specific impact on health resource allocation or the ability of NGOs and non-profits to address the health needs of the healthcare system in which potential subjects participate. Insofar as researchers or sponsors have little reason to believe that the fruits of research – whether knowledge or

²² Shah, Wolitz, & Emanuel, *op. cit.* note 15.

intervention – will have an impact on decisions affecting subjects’ local health systems, this should be considered reason not to include those subjects in a trial.

Before continuing to address what is needed to fully implement the framework, I want to pause to highlight that as it stands, the framework can be useful when applied as a threshold condition. Although it may be difficult to identify what constitutes *sufficient* value to conclude that a trial is worth the costs in human and physical resources and opportunity costs to a community, we can specify a heuristic test of the following form: If the most valuable likely output of a trial (in most cases, rejection of the null hypothesis) would not alter any upcoming policy decision affecting a community, the trial is unlikely to generate locally socially valuable knowledge and should not be conducted in that population. Even this under-specified, threshold condition is sufficient to rule out trials such as Surfaxin that have no immediate potential to improve local healthcare via their outputs.

That said, the SVK Framework if fully specified *is* also able to reflect the fact that high-cost interventions eventually come off patent, and may in the long run become affordable enough for local use. But such long-run benefits are highly uncertain: First, although it may seem obvious that drug prices decline with patent expiration, the opposite has occurred in notable instances. Prices of recombinant factors used to treat hemophilia, for example, have spiraled upward since first introduced in the 1990s, as purer products have been developed and manufacturers have ceased production of lower-purity products that may have eventually become affordable enough for use in LMICs.²³ Second, some interventions remain cost-prohibitive for reasons other than patent protection, such as infrastructural requirements for

²³ J. Kimmelman. Clinical Trials and Scid Row: The Ethics of Phase I Trials in the Developing World. *Dev World Bioeth* 2007; 7: 128-135.

delivery or high manufacturing costs. Finally, we cannot know what the surrounding decision-space facing communities will look like in the future. Each of these considerations introduces significant uncertainty with respect to impacts on future decision-making, and expected value must be discounted appropriately. We should therefore be extremely wary of relying on long-term benefits of this form to justify trials that have failed the heuristic test as applied to relatively short-term decision-making. At a minimum, such long-term benefits must be sufficient to outweigh potential costs to communities in terms of opportunity costs and physical and human resources, even after appropriately discounting for uncertainty. This suggests the bar will be lower for less costly or socially invasive research (e.g. epidemiological studies), but potentially very high for clinical trials that are more invasive or resource intensive.

Because the SVK Framework uses the set of decisions that might be influenced by research outcomes as a surrogate measure of how society values different information, both versions of the framework require some knowledge of the decisions likely to be faced by relevant stakeholders, including policy-makers, funders, and others working locally to address health deficits. Different populations will undoubtedly vary tremendously with regard to the values emphasized in social decision-making, which seems to imply that the framework requires detailed knowledge of the social utility function(s) informing local policy decisions.

In fact, however, such a detailed understanding of how local social values are assessed, prioritized, and pursued in social decision-making is not necessary for external parties to assess a trial for responsiveness. Rather, that the heuristic version of the framework relies on the likely impact of trial outcomes on short-term social decision-making means community priorities can dictate the relevance of research through an understanding of *whether* any social decision is likely to be altered as a result of new information. Those applying the Framework to assess

externally-funded research can determine whether a study is relevant to local policy decisions through open dialogue with policy-makers, regulators, and other organizations before the initiation of a trial. In this way, the SVK Framework is superior to objective health impact measures which may require accurate epidemiological data and appear to deprive local communities of a voice in determining which health deficits will be prioritized.

A critic might object that this approach to defining what is socially valuable conflates what is actually responsive to local health deficits with the interests or priorities of local policy-makers.²⁴ The critic would be right, in the following sense: Insofar as there are no local stakeholders who are willing and able to act on the information a trial might generate, the knowledge will not be put to use in the community, and so cannot provide instrumental social value. Does this leave open the possibility that the account of responsiveness on offer will reflect the decisions that policy-makers are focused on, rather than the decisions that policy makers *ought*, in some sense, to be focused on? It does, although this may be seen to be less problematic by the following considerations. First, the SVK framework is not addressed *only* to policy-makers, whose decision-making might be subject to dubious influence, but also to other local stakeholders, such as NGOs or aid groups working locally. So, for example, during the recent Ebola outbreak in Western Africa, Médecins Sans Frontières was extremely active in the region, and if knowledge from an Ebola vaccine trial could have been useful to MSF in their local efforts to combat the outbreak, this would have been sufficient to meet the criteria, even if local policy-makers were for some reason unwilling to consider said trial in local decision-making.

²⁴ I'm grateful to an anonymous reviewer for pressing these points.

Additionally, however, we should bear in mind that many competing considerations go into health policy decision-making at the local, state, and national levels. Ideally, official decision-making at each is to some extent influenced by popular will via political processes. Obviously this will be true to a greater or lesser extent depending on locale, but insofar as we are concerned with respecting the autonomy of communities, we should want our conception of value to reflect what is valuable *to the communities in question*, not what is valuable on some externally-imposed measure. The SVK framework does exactly this: it indexes social value to the instrumental use to which knowledge will be put in serving local communities, rather than to the instrumental use to which we, as outsiders, would like knowledge to be put in those communities.

Finally, the SVK framework successfully addresses the limitations of reasonable availability. Because it encompasses not only policy decisions regarding investment in healthcare infrastructure and resource allocation, but also what further research to prioritize, the framework can account for the redundancy that is essential to clinical research and can assign value to research outputs that are not immediately related to the production of new effective interventions. For example, the framework can provide assessment of epidemiological research by looking to whether decisions about local health infrastructure are likely to be informed by data on how and whether patients experiencing illness seek care within their communities. Early-phase research can similarly be accounted for: local researchers may be facing decisions regarding how to prioritize limited research resources in combatting locally endemic illness, and early studies can help inform such decisions. On the other hand, local communities may lack the resources to introduce new approaches to healthcare provision in response to epidemiological data, or have insufficient research infrastructure to carry on a fruitful line of research informed

by early phase studies. In such cases, although the information gathered may be valuable in itself, absent local stakeholders who can make use of the information, it fails to provide social value – that is, instrumental value that devolves to the community. Given this, the framework can be sensibly applied to all phases of research, rather than just the late-phase interventional studies to which reasonable availability is most relevant.

But consider the following: trials often generate information along multiple dimensions. One could argue that the most ethical trials are those designed to assess the greatest number of useful hypotheses (thus generating the most social value) while putting the fewest number of participants at risk. This might require the assessment of multiple potential trial outputs. For example, suppose a large study could answer an additional question with a 20% increase in enrollment, whereas answering that question separately would require a second trial as large as the first. Now suppose the second question is less relevant, or of no relevance, to local decision-making. Would increasing the trial's enrollment to answer the additional question be permissible? And if so, subject to what constraints? The answer is unclear, but it does seem like any plausible framework should be able to accommodate some level of such trade-offs. The answer might require more thorough knowledge of the actual value the host community places on the potential knowledge of relevance to them, in the form of the specific expected utilities associated with various outcomes.

Conclusion

In the preceding, I have aimed to do two things. First, I have introduced a framework for assessing the social value of research. Given the social investments in terms of physical and human capital as well as opportunity costs associated with biomedical research, there may well be strong instrumental considerations in favor of a social value requirement. The framework

proposed here gives content to that requirement in a way that is substantive enough to provide assessments of social value and flexible enough to be leveraged to provide insight on a number of different salient questions in contemporary research ethics.

I have also sought to illustrate how this framework can be used to inform the responsiveness requirement in a way that is consonant with the requirement's purported goals. Although there is room for greater precision in applying this framework, its heuristic value should be clear. Given the most valuable potential epistemic output of a clinical trial, if that output would not alter the outcomes of decisions affecting a community's health system in the relatively short term, it is unlikely to generate locally socially valuable knowledge and there should be a presumption against its conduct within that population. Although with further development the framework could function to inform health research prioritization, even this need not entail a strict ordering of social preferences. But as it stands, even without more detailed knowledge of a community's utility functions, the framework provides a useful tool to rule out trials which clearly fail to meet the criterion.

Finally, I should be clear that there are limits to this conception of responsiveness. First, because the framework requires that knowledge gained from a trial have the potential to alter the outcome of some local decision-making procedure, in some instances it will not assign social value to individual trials that could contribute to changing perspectives but fail to do so independently of other research. Given the preponderance of evidence often needed to sway decision-making or inform licensing applications, this could represent a significant shortcoming. However it need not be: the need for multiple trials to inform decision-making might suggest that the framework can be applied more effectively to research portfolios rather than to individual trials. Such an approach may also contribute to reducing the extent to which the SVK

framework prioritizes short-term knowledge gains over longer-term, but potentially more fruitful, research endeavors. A portfolio-based account of responsiveness would necessarily be more complex, however, and more work needs to be done in exploring the feasibility of such an approach.

Second, although the framework can inform decisions about what research it is appropriate to conduct in what settings, and in this way remove certain incentives that might function to preserve or even exacerbate health deficits in LMICs, it does not necessarily provide positive incentives to address the 10/90 gap. For instance, the framework does nothing to prohibit continued research in LMICs addressing low-impact, low-hanging fruit conditions which make up only a small portion of local disease burdens. Such research might inform local policy decisions, depending on how communities prioritize their healthcare and infrastructure spending, but leave the largest needs unaddressed. This simply points to the natural limit of the framework in highlighting one necessary criterion for research conducted in LMICs, but not providing a sufficient condition for ethical research. It is still the case that more ways are needed to positively incentivize that research which is of the greatest social value to the neediest communities.

Acknowledgments

The author would like to thank Derrick F. Gray, Kevin Zollman, Annette Rid, Seema Shah, and two anonymous reviewers at Bioethics for their valuable feedback on earlier drafts of this manuscript.