

An Economic Analysis of Human Subjects Research Ethics

Characterizing the Subjects Rights - Social Benefits Tradeoff

Carl V. Phillips, MPP PhD

 Epiphi Consulting Group, www.epiphi.com

and Minnesota Center for Philosophy of Science, University of Minnesota

phillips@ephiphi.com

Poster #754, Congress of Epidemiology, Toronto
15 June 2001

Based on the working paper: "An Economic Analysis of Human Subjects Research Ethics; Characterizing the Subjects Rights - Social Benefits Tradeoff," Robert Wood Johnson Foundation Scholars in Health Policy Research Program Working Paper WP-17, September 2000. (Not yet submitted. Comments welcome.)

Introduction

Basic principles of economics (the science of trading off scarce resources):

It is almost always too costly -- in terms of whatever else you would have to give up -- to achieve perfection in anything.

Every decision -- including those involving ethics -- requires tradeoffs and it is necessary to have some method for trading off disparate goods and bads.

Some version of this point can be found in most (though not all) studied overviews of bioethics and human subject research (e.g., Beauchamp and Childress, 1994, pp. 28-37; Bulger, Heitman, and Reiser, 1993, p. 143). Typically it is alluded to in most of the literature and the institutional guidelines, but little attempt is made to operationalize it.

Economic Concept of the *Objective function*

Mapping of the disparate actions and outcomes from a particular option, action, or state of the world to a scalar value (i.e., give it a score in the form of a single number) that a decision maker wants to maximize.

Key underlying realities:

- There is no escaping an implicit conversion to a scalar value because it is impossible to maximize more than a single dimension at a time.
- Tradeoffs will be made, whatever our normative ethics or disparate values, and careful considerations of what we are paying and what we are getting for it can help clarify our available options, no matter how politically or ethically charged the ultimate decision is.
- Not only *can* we compare apples and oranges, but we inevitably *do* compare them, at least implicitly, and *should* compare them *explicitly* if we expect to make good honest decisions.

Central Ethical Tension in Health Research

What is best for society is generally different from what is best for an individual study subject.

Deontological (clinician obligation- or patient rights-based) ethic.

Simple version: Every presenting patient (and thus, study subject) should be given the best available treatment (possibly conditional being willing and able to pay for it).

-versus-

Consequentialist social welfare goal.

Net social benefit, largely in the form of better information about how to treat future patients, can justify a net expected¹ loss to the presenting subject/patient.

¹*Net* refers to subtracting all costs of an action (or situation, etc.) from the sum of all benefits. *Expected* refers to taking probability-weighted averages across all probabilistic outcomes, based on our ex ante beliefs about the probabilities.

The Need to Recognize and Analyze the Tradeoff

We, as a society, have decided that we are going to partially fulfil the objective of caring for the individual patient/subject and partially fulfil the objective of improving social knowledge and thus the care of people in the future.

We need to carefully identify the available options.

We need to think about where we want to draw the line (rather than claiming that we are not doing so).

Researchers very rarely acknowledge "**this research harms the study subjects in identifiable ways, but the benefits outweigh the costs.**" They often insist that the experiment does not really harm subjects, which cannot be true (see aside, below).

Even if they acknowledge it, **they almost never make any effort to assess the expected benefits**, so the claim is baseless. There is virtually no discussion of how to quantify the expected social benefits of a health research study ex ante, even though it is possible to do so (Phillips, 2001, in press).

Aside: Human Subjects Experiments Always Harm Study Subjects -- the simple logic from economics and decision theory

Common statements of the ethical obligation to best treatment often impose either *impossible constraints* on controlled trials or virtually *no constraint*.

No constraint: precluding only the assignment of treatments that are *known* to be worse (as in the claim, "we do not *know* which treatment is better until we complete the study"). Since it is impossible to be 100% certain about the causal relationship between a treatment and outcome -- we cannot know the counterfactual outcome from the action not taken, even *ex post*, let alone *ex ante* -- this standard will never be violated.

Impossible constraint: precluding a treatment other than the "best" available.

What does it mean to give the "best" treatment? Any definition of "best" that can be applied by non-omniscient decision-makers requires the caveats, "based on currently available information" (with the recognition that forthcoming information might lead to a different assessment) and "more likely than not for this particular individual given her characteristics" (with the recognition that there are uncertainties and the therapy might actually be worse for her, but we can only play the odds). Thus, there will always be a best treatment for any presenting patient. Thus, assignment that is at all influenced by study needs violates the constraint.

The most common method for alluding this is disingenuous, the appeal to "equipoise" -- that a particular clinician or other agent might be totally unsure about which of two particular treatment options is more promising for a particular individual. After adequate reflection and study, the chance that a clinician truly cannot identify one treatment as likely to be better than another on average approaches zero (Freedman, 1987). It is almost impossible that this indifference would independently apply to every single trial-eligible patient who presents to him (Appelbaum et al., 1987). Claims like "we have not completed our trial, so I am not sure the treatment I believe is superior on average really is, and therefore I am not giving anyone an inferior treatment when I assign her treatment at random," demonstrate a (possibly willful) dismissal of the fact that there is always a best decision based on current information. After all, once the research is finished, decisions will be made based on a new, but still imperfect, collection of currently available information.

Official Ethical Guidelines Offer Little to Help us Make the Tradeoff between Deontological and Consequentialist

Advice on making the tradeoff often uses purely consequentialist language.

Declaration of Helsinki: "research involving human subjects cannot legitimately be carried out unless the importance of the objective is in proportion to the inherent risk to the subject," and "should be preceded by careful assessment of predictable risks in comparison with foreseeable benefits to the subject or to others" (World Medical Association, 1964).

Nuremberg Declaration: "The degree of risk to be taken should never exceed that determined by the humanitarian importance of the problem to be solved by the experiment" (*Trials of War Criminals*, 1949).

Belmont Report (National Commission for the Protection of Human Subjects, 1979), which dominates U.S. government regulations: discusses obligations to patients as a principle of beneficence, but in the operationalizing "Applications" section (and the U.S. government and research institution policies that adhere closely to it) rights and obligations are reduced to informed consent and the consequentialist risk-benefit comparison.

Or strong deontological interpretations effectively preclude any experimentation.

Authors suggest that all human experimentation is precluded (classic reference is Fried (1974)).

Self-contradicting the above, Declaration of Helsinki: "Concern for the interest of the subject must always prevail over the interests of science and society."

(There has been surprisingly little assessment of whether these central works on the subject improve study ethics other than by truncating the most egregious cases. They certainly do not much help us with the unavoidable need to balance competing ethical goals, given that our actions show we want to do so.)

Economic Concept of the *Possibility Frontier*

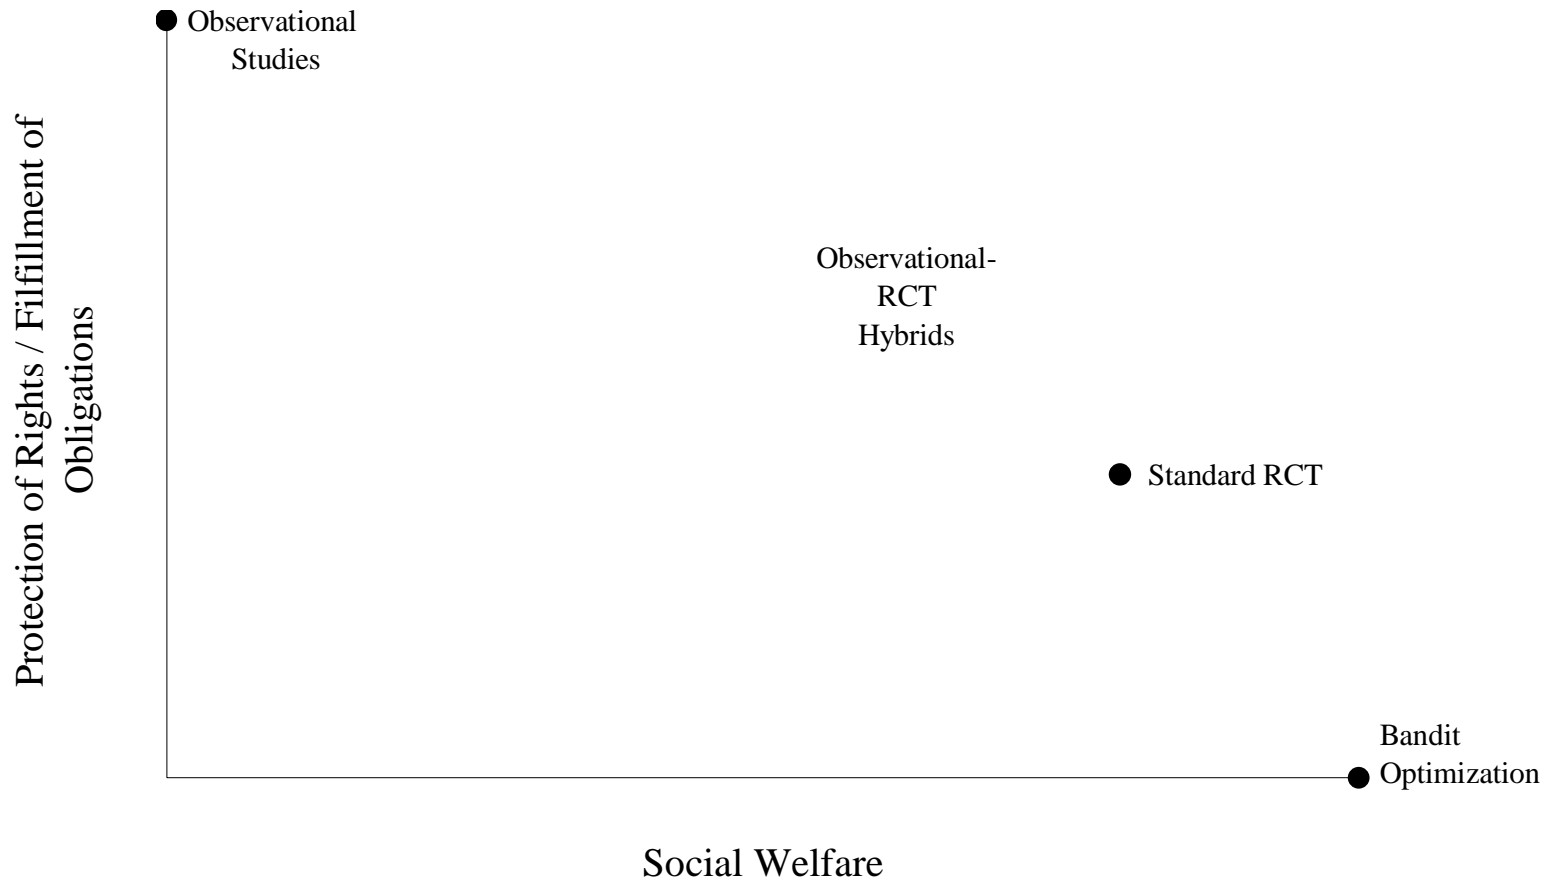
An n-dimensional function describing "physically" possible tradeoff among different "goods" (goods = anything people want).

Can be thought of as the maximal production of one good given some minimal production of the other(s).

E.g., trading off leisure time vs. wages; guns vs. butter; or deontological vs. consequentialist goals in a study.

Figure 1 simplifies the research tradeoff to two dimensions (ignoring dollar costs, in particular) and notes a few points among the many that form the frontier. (Possibility frontiers are typically drawn as continuous, but they need not be.)

Figure 1
Possible Choices for the Ethical Tradeoff



Y-axis, Social Consequences: long-run welfare of all members of society (including study subjects) resulting from our action (a particular research strategy).

X-axis, Clinical Obligation/Patients' Rights: degree of adherence to deontological rules of patient rights or clinician duty to provide best treatment

Scales are necessarily conceptual rather than precise quantities (though bounds take natural values, as discussed below).

Measure of rights/duties very rough, and possible best thought of as ordinal.

Consequentialist social welfare function can be quantified, but this presentation glosses over choice of health measures and how to aggregate.²

² (Note that in virtually all discussion, the ostensible measures of social benefit are not presented in terms of actual social benefit achieved for the cost. Instead, the social welfare side of the research tradeoff is nearly always represented in terms of some test of statistical power, which is clearly not the right measure (see also, Phillips 2001).)

Immediate Conclusions resulting from Possibility Frontier Concept

It is difficult to agree on tradeoffs, but we can always formally describe the tradeoff options, potentially clarifying the debate.

Even if we do not agree on the shadow price, we ought to agree that we would like to get to yet unreachable points above the curve (by improving our technology), and that those options below the curve are inferior.

Tradeoffs are a willingness to trade one good for the other at a particular exchange rate, the "*shadow price*."

We make decisions based on implicit shadow prices, whether or not we carefully consider them. With each research study, we are implicitly stating "social benefit X_1 is worth cost to subjects Y_1 , but the additional benefit X_2 is not worth additional cost Y_2 ."

Often, the implicit prices contradict each other and are contrary to the explicit price we would choose, as becomes obvious with this analysis.

Obscuring the tradeoffs because we do not want to think hard about them (while making them anyway) is a recipe for bad policy.

Discussions of this general point can be found in many policy arenas, including consumer safety and worker health (e.g., Viscusi, 1998), environmental law and policy (e.g., Phillips and Zeckhauser, 1995), medical care (e.g., Gold et al., 1996), and general public policy analysis (Stokey and Zeckhauser, 1978).

Putting Research Options in the Possibility Frontier Context

Observational Studies. The only way of avoiding rights/duty violation.

Minimizes the burden on the subjects, and thus provides the maximal value for the y-axis (and the zero value for the x-axis is set based on this).

Sometimes the only practical option, but also used to study exposures when a trial is practical.

Will not interfere with assignment of best treatments, but will produce lower expected social welfare than a comparably sized study that allows the researchers to assign treatments/exposures (due to extra confounding, etc.). Thus, all other study designs appear to the right of the axis.

Randomized Controlled Trial (RCT).

Gives some subjects an inferior treatment with certainty, putting the method lower on the y-axis than observational inference.

Its popularity demonstrates that we are willing to yield on clinical obligation / patients' rights to improve welfare.

RCT that is comparable in size (statistical power) should give superior data compared to the observational study. Thus it will improve (in expected value terms) the social welfare by increasing the chance that we find the right answer.

Multi-Armed Bandit Optimization. Social-welfare optimizing research strategy.

It is always possible to construct a method to maximize a single well-defined objective function. In this case, it is the solution to the class of optimization problems is known as *multi-armed bandit* problems (a reference to the slot machine, the "one-armed bandit") or Gittins optimizations (Gittins, 1989).

If a slot machine has two levers (or we have a choice of two health-affecting actions) and you are uncertain about their actual payoffs, there is a tension between wanting to keep using the one that apparently gives the better payoff based on current information and using the other to learn more about its actual payoff (which might be higher than you currently believe). To optimize social welfare we need to gather some optimal amount of data about a treatment we think is inferior before giving ourselves over to the other for every future patient. It is possible to solve this optimization numerically and thus generate maximal social welfare.

Maximizes social welfare, and thus provides the maximal value for the x-axis (and the zero value for the y-axis is set based on this).

In particular, it produces higher social welfare than a RCT, which is not likely to assign the right number of subjects (in particular, is likely to stop too soon).

Guarantees researchers will face a glaring ethical problem. The study will reach a point where the researchers are very close to certain that one treatment is superior, but not sure enough to stop the experiment (after all, if the treatment of a million future patients is at stake, it might be socially superior to give one more person the treatment that appears more than 99% certain to be inferior, just to make sure it really does not work well).

Presumably as a result of this, serious consideration is not generally given to designing research in this way, though it may be to lack of knowledge of the option.

Alternative Research Designs

Discussions of research alternatives typically centers around the choice between observational inference and the RCT. Options to the right of the RCT are seldom discussed. Even options in the space between observation and the RCT are seldom explored.

Unbalanced RCT. Potential improvement on both dimensions.

An RCT can usually be improved by having very different numbers of subjects in each treatment arm. When there is already a stable estimate of the effect of one treatment option (e.g., the existing standard practice), control data is only needed to confirm that the outcome for the study population is as expected.

If we believe the new treatment is superior, ethical advantage of assigning fewer subjects to the inferior treatment (to the right of RCT). Still violates an absolutist interpretation of obligations/rights, so below observational studies.

For a given study size, more subjects can get new treatment, increasing information from study (higher than RCT).

If done badly, may be worse than RCT on either or both dimensions.

Need to make optimal use of prior knowledge, requiring non-standard statistical methods (and loses epistemic advantage of RCT in dealing with counterfactual problem).

RCT with a Stopping Rule. May be worse along both dimensions.

Goal of stopping a study early is obligation/rights ethics: to reduce the amount of person-time that subjects are assigned inferior option. However, it is actually ethically ambiguous.

Preliminary review of data can require researchers to assign a treatment that they have become very sure is inferior, but not enough to meet stopping rule. This might be considered worse than when there is a weaker belief.

Could provide a limited version of the consequentialist advantages of a bandit optimization, but as currently practiced *cause net harm*. Stopping points have much less certainty than is optimal (e.g., reducing by another 1% the chance of given the wrong advice to millions of people is likely worth continuing the study).

Observation-Trial Hybrids. Easing the violations of RCTs while minimizing confounding.

One set of hybrid strategies randomizes patients to different clinics which each use their preferred treatment (Korn and Baumrind, 1991). More nuanced versions involve intermediate updating of preferences by each clinic, thereby rebalancing the study arms (Kane, 1998; Phillips and Kane, unpublished).

Other strategies assign treatments based on cutpoints of objectively quantifiable subject characteristics (Trochim and Cappelleri, 1992; Cappelleri and Trochim, 1994, Campbell and Stanley, 1966). The presumably better treatment can be given to patients who can benefit more, and statistical assumptions can correct for the lack of a direct comparison of treatment groups.

Treatment assignment is not based on unobserved clinical evaluation of subject characteristics, reducing some major sources of confounding, so between RCT and observational studies in information value.

Minimizes, but does not eliminate, clinicians administering what they believe are inferior treatments to needy patients, so better than RCT for duty/rights, but worse than observational study. (Begs the question of why we find it acceptable for different clinicians or groups of clinicians to recommend and use different treatments.)

Reduces the incentive for clinicians to introduce bias (likely unrecognized by those analyzing the data) by cheating in the treatment assignments to make sure that certain potential subjects get the treatment that the clinician prefers for them, regardless of random assignment (Schulz, 1995).

A clever design, taking advantage of the specifics of a situation, could get close to the best of both worlds, pushing the possibility frontier up and right.

Summary Conclusions

Any experimentation will result in someone's health being diminished. Researchers have an obligation to ensure that their expected benefits to society justify this cost. This requires understanding and measuring what the social benefits are.

The possibility frontier approach helps clarify and organize the spectrum of choices and available tradeoffs.

We have ethical obligations to not simply take study designs off the shelf. The possibilities for choosing, combining, and tweaking study designs can offer a wide menu of practical and ethical impacts. Part of every research plan should be to consider the implication of these options. Since many implications are not easily predicted without quantitative analysis, modeling and simulating the results of the research is likely to be fruitful. It is not enough to pick a method that is theoretically better in certain ways, because the rules-of-thumb might be wrong.

REFERENCES

- Appelbaum, PS, LH Roth, CW Lidz, P Benson, and W Winslade (1987).
- Beauchamp, TL and JF Childress (1994). *Principles of Biomedical Ethics, 4th edition*. New York: Oxford University Press.
- Bulger, RE, E Heitman, and SJ Reiser, eds. (1993). *The Ethical Dimensions of the Biological Sciences*. Cambridge: Cambridge University Press.
- Campbell DT and JC Stanley (1966). *Experimental and quasi-experimental designs for research*. Chicago: Rand McNally.
- Capperlari, JC and WMK Trochim (1994). An Illustrative Statistical Analysis of Cutoff-Based Randomized Clinical Trials. *Journal of Clinical Epidemiology*, 47(3):261-270.
- Freedman, B. Equipoise and the Ethics of Clinical Research. *The New England Journal of Medicine*, 317(3):141-145.
- Fried, C (1974). *Medical Experimentation*. New York: American Elsevier.
- Gittins, JC (1989). *Multi-armed Bandit Allocation Indices*. Chinchester: John Wiley & Sons.
- Gold, MR, JE Siegel, LB Russell, and MC Weinstein (1996). *Cost-Effectiveness in Health and Medicine*. New York: Oxford University Press.
- Kane, D (1998). *Disagreement*. Doctoral dissertation, Harvard University
- Korn, EL and S Baumrind (1991). Randomised clinical trials with clinician-preferred treatment. *The Lancet*, 337:149-152.
- The National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research (1979). Ethical Principles & Guidelines for Research Involving Human Subjects (The Belmont Report). U.S. Department of Health, Education, and Welfare, 18 April 1979. (<http://ohrp.osophs.dhhs.gov/humansubjects/guidance/belmont.htm>)
- Phillips, CV (2001, in press). The Economics of "More Research is Needed". Forthcoming in *International Journal of Epidemiology*.
- Phillips, CV and Kane D (draft). Decentralized Trials as an Alternative to RCTs.
- Phillips, CV and RJ Zeckhauser (1995). Evaluating the Present Natural Resource Damages Regime: The Economist's Perspective. In *Natural Resource Damages; an appraisal of the regulatory and liability regime*, Richard B. Stewart, ed., New York: National Legal Center for the Public Interest.
- Schulz, KF (1995). Subverting Randomization in Controlled Trials. *JAMA*, 274(18):1456-1458.
- Stokey, E and RJ Zeckhauser (1978). *A Primer for Policy Analysis*. New York: W.W. Norton.
- Trials of War Criminals before the Nuremberg Military Tribunals under Control Council Law No. 10, Vol. 2*. Directives for Human Experimentation (Nuremberg Code). Pp. 181-182. U.S. Government Printing Office, 1949. (<http://ohsr.od.nih.gov/nuremberg.php3>)
- Trochim, WMK and JC Capperlari (1992). Cutoff Assignment Strategies for Enhancing Randomized Clinical Trials. *Controlled Clinical Trials*, 13:190-212.
- Viscusi, WK (1998). *Rational Risk Policy*. Oxford: Clarendon Press.
- World Medical Association (1964, amended 1975, 1983, 1989). The Declaration of Helsinki: Recommendations Guiding Medical Doctors in Biomedical Research Involving Human Subjects. Adopted by the 18th World Medical Assembly, Helsinki, Finland, June 1964. (<http://ohsr.od.nih.gov/helsinki.php3>).