# Experimental Economics' Ethics

By: John Ifcher and Homa Zarghamee

On September 18, 2010, the New York Times featured a piece on two cousins, Thomas McLaughlin and Brandon Ryan, both in their early twenties, both battling the lethal skin cancer melanoma (Harmon, 2010). Thomas's initial diagnosis was much worse than his cousin's—so much so, that he was eligible to participate in a clinical trial for a new drug, PLX4032. He was assigned to the treatment group and saw miraculous rates of improvement; meanwhile, Brandon's health deteriorated with the progression of the melanoma and under the phenomenal physical stresses of chemotherapy. It therefore came as a mixed blessing when Brandon finally became sick enough to qualify for the next round of PLX4032's clinical trials. Any potential upside of this mixed blessing, however, was quickly dispatched upon news that he had been assigned not to the treatment group receiving PLX4032, but to the control group—the counterfactual, the group that the miracle-drug's rates of success would be compared to and in which his ineffectual chemotherapy would continue.

To maintain experimental tidiness (and, importantly, validity), assignment to the control group in clinical trials precludes any alternate routes of access to the drug being tested and does not allow patients to participate in future trials if they drop out. Much to the chagrin of Brandon, his family, and his doctors, Brandon was stuck in the control group with absolutely no access to the treatment, the treatment that was working wonders on his cousin. So successful was PLX4032 in its trials and so mild were its side effects (compared to chemotherapy) that even the most assuredly learned and devout followers of the experimental method—that is, much of the medical-research community itself—believed it unethical to block the control group from getting PLX4032 until the experiment's end. A raging ethical debate ensued in the research community that boiled down to whether the hard, scientific knowledge derived from completing the trials was, indeed, the greater good. In the end, Brandon succumbed to his illness and died at age 22. True, Thomas's fate would likely have been the same without the PLX4032 trials, but the fact remains that in the ethical calculus of such trials, for every Thomas there had to be a Brandon.

Compared to what medical researchers face, the ethical considerations of experimental economists may seem like child's play. At least a few experimentalists we know admit that, when first confronted with the pages-long application for approval to use human subjects in their research, they thought of it as an annoyance so pointless and mechanical that it must have been a bureaucratic mistake that their university required them to go through the process. Surely, they thought, such approval was really meant for researchers like the ones testing PLX4032; surely, sooner or later the university Institutional Review Board (IRB) would notice that experimental economists weren't working on issues of life and death and spare them the paperwork.

But to deny the importance of experimental economists' ethical considerations on the grounds of seemingly distant, if any, relation to issues of life and death is to deny the

importance of economics research on the same grounds. Economic experiments—implemented from sterile university laboratories to the ministries of health and education in developing countries—test the theories and evaluate the social policies that shape the economic environment in which people live and the resulting choice-sets that they confront. The findings can have profound impacts. For example, it was experimentally demonstrated that changing from an " opt-in" to an "opt-out" 401(k) pension program increased employee participation from about one-third to over 80 percent (Madrian and Shea, 2001). Johnson and Goldstein (2003) studied organ donations programs, where "opt-in" versus "opt-out" defaults have a profound impact on organ donation participation rates. And of course, any decent economist knows to, at a bare minimum, speculate upon the policy implications of her work; "why does it matter?"

The answer to this question is on the benefits side of the cost-benefit analysis of doing research. Allowing for a little romance, maybe it *alone* drives the researcher to spend her days, months, years deriving the proofs of her theorem, obtaining just the right data set with just the right measures and analyzing it from every possible angle. To the experimentalist studying the impact of X on Y, the right data set with the right measures may have to be one designed *with the research question in mind*. That is, the effect of interest may be impossible to tease out of observational data because of the difficulty of proving that the people to whom X happened and did not happen were otherwise identical. So she makes certain that the subjects assigned to X and not-X are assigned randomly and not on the basis of any observed or unobserved characteristics: this is critical to experimentalists' identification strategy (for demonstrating a causal link) and knowledge-generation, and its violation undermines the methodological benefits of the experimental method. The key point is that the experimenter conducts her research by *exerting control* over her subjects' environment and *manipulating* their choice sets. So it is incumbent upon her to add to the cost side of her cost-benefit analysis the effects of her control and manipulation and how these effects weigh against the answer to "why does it matter?"…if not naturally so, then certainly once the IRB-cum-"benevolent social planner" has made her internalize the externalities of her research method.

Below, we offer our thoughts on the ethics of economic experiments. Because we are economists and somewhere along the way our brains' hemispheric division became cost-benefit instead of left-right, we will continue to fall back on that framework in hashing out the impact of an experimenter's choices. We don't claim to be the originators of the many of these thoughts. Despite being the last social science to adopt a code of ethics and despite this code of ethics having nothing to say about experimentation or human subjects (or really anything other than citing sources of funding and being up front about financial conflicts of interest in published work), a number of economists have written about the ethics of experiments, and our insights draw from their writing in addition to our own experiences. As in any cost-benefit analysis without common units of measurement, ours will not lead to decisive rules on how to run experiments, but hopefully it will provide food for the introspective experimenter's thought. A final obvious but important note: in the words of Karen A. Hegtvedt, "researchers do not have an inalienable right to pursue research with human subjects (Hegtvedt, 2007, p. 159)" My wanting to know the economic consequences of war—an undoubtedly important

economic question—does not give me a carte blanche to start one. I may have to resign myself to accepting that, like a world without World War II, the proper experimental counterfactual is untenable. If a question cannot be answered without breaches of ethics, then the experiment shouldn't be conducted.

Let's first consider some of the actors affected by an economic experiment. The researcher (either with or without a non-research partner or sponsor) uses human subjects (either with or without their knowledge) to measure the causal effect of X on Y. The researcher herself stands to gain from this research, as do the targeted population of the research (e.g., any individual or institution trying to change its Y) and other researchers working on related themes (e.g., researchers studying the effects of X, researchers identifying the determinants of Y, researchers using experimental methods, etc.). The experimenter's research choices will impact these different actors in obvious and not-so-obvious ways. We will first consider the ethics of using human subjects and then the ethics of reporting and publishing results.

## Ethics of Using Human Subjects

### *Ethical Guidelines*

The most obvious group affected by experimentation is the pool of human subjects. In the lion's share of economics experiments, there are no real threats of the physical suffering and side effects that may be inflicted by biomedical experiments, nor is the psychological stress ever going to compare to what Brandon Ryan went through being denied PLX4032. That said, the direct benefits to the human subjects of economics experiments will never come close to those enjoyed by Brandon's cousin Thomas. The potential costs and benefits to human subjects in economics are usually much more subtle. Perhaps because of the subtleties, it is important to know the ethical standards used by IRBs to approve the use of human subjects. The common source of these standards is the Belmont Report, created in 1978 by the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. The Belmont Report puts forth three guiding principles: respect for persons, beneficence, and justice. Respect for persons asks that the autonomy of individuals be respected, that they give informed consent to participate in the experiment, and that individuals with diminished autonomy be protected. Beneficence stresses the Hippocratic maxim to "do no harm"—that is, to maximize the benefits and minimize harm to human subjects. Justice requires that subjects be chosen non-exploitatively and that the benefits of the research be available to those burdened by it.

In the infamous Tuskegee syphilis study, the untreated syphilis-progression of 399 impoverished African American men was tracked for 40 years and compared to the health of 201 non-syphilitic men. The subjects, made to believe that they were receiving free government health care, were never told they were part of the study, that they had syphilis, or that they could be cured with penicillin. In fact, there are reports that the researchers blocked access to alternative treatments when they became available. It

should come as no surprise that Tuskegee fails to meet any of the Belmont Report's three guiding principles…that, in fact, the Belmont Report was created in response to the whistle blown on Tuskegee. What may come as more of a surprise is that another controversial set of experiments, Stanley Milgram's Yale experimentes (1963, 1965, 1974), does not categorically fail to follow the three guiding principles. Suspecting that the magnitude of Nazi violence was at least partially rooted in obedience to authority, Milgram tested the willingness of subjects to administer increasingly powerful electric shocks to confederates masquerading as subjects, whose (albeit fake and scripted) objections, pleas, and screams could be heard from an adjacent room, all under the guise of a laboratory experiment on memory and learning. After a certain voltage, the confederate falls silent, suggesting that he may have passed out or died, and the experimenter urges the subject to continue. Despite having been told up front that they were free to leave the experiment at any time without forfeiting the show-up fee, a whopping 65% of subjects administered shocks up to the maximum voltage of 450 volts.

Milgram took care to debrief his subjects at the end of the experiment, to follow up with them regularly, and to attempt to relieve them of any resultant guilt, lost self-esteem, or identity crisis. To many critics, this was hardly enough, but Milgram's student Alan C. Elms argues that, if anything, it was too much, that Milgram was under no ethical obligation to bolster subjects' false sense of self-esteem (Elms, 1972). In a rare case of real-time experimental education, the vast majority of subjects themselves were actually glad to have participated, many claiming to have learned an important lesson about themselves and human nature. So careful was Milgram that his protocol had to be only lightly tweaked (psychological pre-screening, closer observation for signs of stress, one third the maximum voltage, and more immediate debriefing) by Jerry Burger to be repeated with full university IRB approval in 2009 (Burger, 2009). The real harm of the Milgram study is the psychological distress the subjects feel when they think they are seriously injuring the confederate. The real achievement of the study is showing that the psychological distress doesn't stop the subjects from taking the actions that will seriously injure the confederate. The interesting point is that it isn't the deception in Milgram's study that would make it unacceptable to a contemporary IRB, nor is it that the subjects suffered psychological distress: both of these are present in Burger's follow up. It's just a matter of degree. So even with the Belmont Report and IRBs, we don't get concrete ethical answers, just a framework for ethical thought.

Christopher B. Barrett and Michael R. Carter (2010) apply the framework, with equally ambiguous results, to Gugerty and Kremer (2008), an experimental study of the Rockefeller Effect: "Taking its cue from John D. Rockefeller, who refused to give money to Alcoholics Anonymous on the grounds that the money would undercut the organization's effectiveness, the Gugerty and Kremer (2008) article explicitly sets out to determine whether grants of money to women's organizations in Kenya distorts them and leads to the exclusion of poorer women and their loss of benefits. Donor groups were providing grants to women's organizations on the presumption that they were doing good. Proving otherwise, and that the Rockefeller Effect is real, could of course be argued to bring real social benefit. However, the ethical complexities of undertaking research designed to potentially harm poor women are breathtaking. Standard human

subjects rules require: (1) that any predictable harm be decisively outweighed by social gains; (2) that subjects be fully informed of the risks; and, (3) that compensation be paid to cover any damages incurred. It remains unclear whether these rules were met in the Gugerty and Kremer (2008) study, which is somewhat chilling given that the study indeed confirms that poor women were harmed by the injection of cash into randomly selected women's groups (Barrett and Carter, 2010, p. 520)."

*Deception*

Since the hypothesized, or at least expected, outcome was a counterintuitive negative one, it may even be considered deceptive of the researchers to have withheld information about the Rockefeller Effect from the subjects. That said, the common practice of withholding the hypothesis being tested and the full breadth of the experiment is not generally considered to constitute deception. Rather, deception occurs when experimenters convey false or intentionally misleading information to subjects. The use of deception in economics experiments is essentially forbidden (by virtue of the impossibility of getting deception past journal referees), and, as a matter of course, the discipline's distaste for deception is often the first thing subjects are told in economics experiments.

Deception can benefit the researcher by increasing the range of questions she can answer. Its costs, to economists, are threefold. First, economics experiments are designed around monetarily incentivized decisions. Deception in the context of financial rewards would quite simply constitute fraud. Second, it may exacerbate feelings of objectification in the subjects and call into question their ability to exercise autonomy in the experiment. Third, deception, especially when institutionalized as it has been in psychology, breaks down the potency of the monetary incentive in all experiments by calling its veracity into question or by supplanting it, if only partially, with other incentives—for example, the incentive to outsmart the experimenter or "spot-the-deception." In their 2001 compilation of the existing experimental evidence both for and against the use of deception, Ralph Hertwig and Andreas Ortmann (2001) recount the real anecdote of a subject's epileptic seizure going initially ignored by other subjects because they thought it was an experimental hoax (MacCoun and Kerr, 1987). More broadly, a researcher's choice to employ deception creates a negative externality for other researchers, present or future, who want to conduct behavioral research without it: the external validity of the subject pools' psychological state in the experiment will be reduced and the credibility of their research will be compromised. In the words of Hertwig and Ortmann, "participant's trust is a public good worth investing in (Hertwig and Ortmann, 2001, p.398)."

*Informed Consent and Blindness*

Deception relates closely to informed consent, one of the cornerstones of the Belmont Report; too much deception may render informed consent moot, as the veracity of the information and what exactly subjects are consenting to is called into question. Given the

paucity of deception in economics experiments, satisfying the right to informed consent is usually as simple as obtaining a signature from subjects approving the general nature of the study in which they will be involved and reassuring them of their freedom to abstain from any or all of it if they so choose. The situation becomes more complicated in natural field experiments in which subjects are not made aware of their involvement in a randomized experiment, either because the nature of the intervention is naturally occurring (e.g., manipulating the wording of a political contribution solicitation letter) or because of any of a slew of named "effects" that would cause subjects to change their behavior because they know they are in an experiment. The John Henry Effect occurs when subjects in the control group take actions to overcome the real or perceived disadvantage of their random assignment. Hawthorne Effects occur if subjects in either control or treatment suspect that the experiment's hypothesized results will be used in negative ways and hence modify their behavior to eschew the hypothesized. The Pygmalion Effect occurs when subjects' actions and perceptions respond not necessarily to the treatment itself, but to meet the expectations of the treatment's hypothesized effect. The likelihood of these possible effects—which are quite real given that human subjects, unlike plots of soil, are active agents—and the feasibility of addressing them with blinded studies must be weighed against violations of informed consent.

The "effects" described above can occur when subjects are not blind to their own treatment versus control status. A related set of concerns can arise when the experimenter is not blind to subjects' treatment versus control status. As Gary Burtless notes, "Except among philosophers and research scientists, random assignment is often thought to be an unethical way to ration public resources (Burtless, 1995, p. 74)." Assuming that the experimenter herself—like, for example, Brandon Ryan's doctors—can, in the name of maintaining experimental validity, stomach the difficulties of withholding treatment from deserving or distressed members of the control group, policy partners and implementers may not. As Barrett and Carter note, "in our experience the unfairness and wastefulness implied by strict randomization in social experiments often sows the seeds of some implementers' breach of research design. Field partners less concerned with statistical purity than with practical development impact commonly deem it unethical to deny a control group the benefits of an intervention strongly believed to have salutary effects, or to knowingly treat one household instead of another when the latter is strongly believed likely to gain and the former not. Well-meaning field implementers thus quietly contravene the experimental design, compromising the internal validity of the research and reintroducing precisely the unobserved heterogeneity that randomization was meant to overcome (Barrett and Carter, 2010, p. 521)." Thus, when feasible double blind experiments are best.

*Monetary Incentivization*

Our earlier discussion of "no deception" brought up another pillar of experimental economics: monetary incentives. Economists are notoriously suspicious of the Bradley Effect (in the 1982 California gubernatorial race, opinion polls favored the black candidate Tom Bradley, and his loss was understood as money not being put where the mouth is); economists think talk is cheap; preferences cannot be credibly spoken of, they

can only be revealed when there are stakes involved. So we have institutionalized the use of monetary incentivization in our experiments to gain the benefit of credibility. One obvious cost is the financial one incurred by the experimenters. Less spoken of are the ethics of monetary incentives. On the one hand, they provide compensation for any inconvenience, boredom, or harm that may arise in the experiment. On the other, monetary incentivization may be seen as an instrument of coercion or exploitation. In order to recruit subjects and inform them of the benefits of participating, experimental economists usually advertise the opportunity to make money, publicizing both the minimum show-up fee, and the average and maximum payout. Despite being ensured of the show-up fee, subjects are often attracted by the extra money (beyond the show-up fee)_they may receive and feel compelled to stay in the experiment or act in accordance with the experimenter's wishes in order to receive it. Remember that Milgram's participants were initially told that they could leave the experiment at any point they wished and still receive payment, but clearly they perceived otherwise later in the experiment. Interestingly, while greater experimental stakes are appealing to economists on the grounds of increased validity, they actually exacerbate the problems of coercion and exploitation. Consider a field experiment in a developing country where the stakes are equivalent to a household's monthly wages. While it can be argued that the high stakes make the subjects' choices that much more real and important, the high stakes may make the subjects feel compelled to do what they think the experimenter wants them to do, or to stay in an experiment against their better judgment instead of settling for the nominal show-up fee.

Another problem with monetary incentives is that they may reward bad behavior. Barrett and Carter report: "As but one prominent example involving widely respected scholars, Marianne Bertrand et al. (2007) randomized incentives for subjects in India who did not yet possess a driver's license, so as to induce them to bribe officials in order to receive a license without having successfully completed the required training and an obligatory driver safety examination. The very predictable consequence of such an experiment is that it imperils innocent non-subjects – let alone the subjects themselves – by putting unsafe drivers on the road illegally. This is irresponsible research design, yet the study was published in one of the profession's most prestigious journals (Barrett and Carter, 2010, p. 519-20)." David T. Dearman and James E. Beard (2009) argue that principal-agent experiments, in which the agent earns more by trumping up her hidden costs, foster deceptive unethical behavior. This is particularly troublesome in university labs, which are, like it or not, embedded in a learning environment and use students as subjects. A related but much more subtle point raised by psychologists is that financial incentives may dampen intrinsic motivations and interest in the choices made in the lab. Protecting subjects from boredom may not be an ethical matter, but it may be of importance to the extent that the change in motivations may diverge with what would be observed of subjects outside of the lab. It may also give rise to a culture of "professional subjects" who use experiments as a regular source of income without putting any real thought into the choices they make while there.

**Ethics of Reporting**

As much as we may want to believe that some mix of intellectual curiosity and wanting to help the world are the only motivation for our research, we all know that, at a practical level, our careers are built on the success of our research—and this success is often measured not by the conscientiousness with which our experiment was conducted, but by the publishability of its results. Ethical scrupulousness in reporting results protects the public and scientific community from fraudulent results and a general loss of trust in research; the researcher should consider a host of questions:

First, has there been any violation of the exogeneity assumption in random assignment? Such a violation could arise, as noted above, if a field implementer did not stick to the random assignment either by accident or for her own reasons. Or we may find out from exit surveys that randomization was not valid ex post. Experimentalists are divided about the proper response. We are personally of the opinion that ex post invalid randomization should be controlled for (Barrett and Carter fall on our side), but we have been advised by senior colleagues at conferences and in referee reports that anything but reliance upon ex ante randomization becomes too subjective; more on this below.

Second, to what extent do the results truly reflect treatment effects or behavioral responses to other features of the experiment? Again, unlike research with passive agents in which the treatment given is the treatment received, human subjects' behavioral responses may be in response to other features of the experiment other than treatment, and which the experimenter is not aware of. For example, a change in the experimental setting.

Third, if sponsored by a partner, to what extent is the experimenter biasing her methodology to obtain results that appeal to the partner? As noted above, the American Economics Association Code of Ethics requires that funding and financial conflicts of interest be reported. Surely, reporting such a conflict of interest is important, but being conscious of it in the design and implementation process is equally important. Glenn Harrison and John List give the example of paired-audit field experiments: "the Urban Institute makes no bones about its view that discrimination is a widespread problem and that paired-audit experiments are a critical way to address it…. There is nothing wrong with this apart from the fact that it is hard to imagine how volunteer auditors would not see things similarly. Indeed, Heckman (1998, p.104) notes 'auditors are sometimes instructed on the problem of discrimination in American society prior to sampling firms, so they may have been coached to find what audit agencies wanted them to find.' The opportunity for unobservable to influence the outcome are potentially rampant in this case (Harrison and List, 2004, p. 1038-9)."

Lastly, are negative results being reported? And are the reported results in response to the original research question or coming from post-hoc analysis? More specifically, negative results are often not as interesting to journals and referees as positive results, so they go unpublished. We have personally conducted costly experiments only to find out later that our colleagues have conducted the same ones years before with negative results. This may incentivize data-mining. For example, consider the analysis of pilot experiments. A

pilot may only be called that ex post when results are not obtained. Or it may be dropped from analysis with the main experiment for no other reason than that it reduces statistical significance. Similarly, justifications may be made for dropping observations that likely would not have been dropped if statistical significance had been achieved in the first shot. Such practices surely confirm the suspicions of those who argue for reliance on ex ante randomization: everything else might just be a trick.

References


Barrett, Christopher B. and Michael R. Carter. 2010. "The Power and Pitfalls of Experiments in Development Economics: Some Non-random Reflections." Applied Economic Perspectives and Policy, 32(4): 515-48.

Bertrand, Marianne, Simeon Djankov, Rema Hanna, and Sendhil Mullainathan. 2007. "Obtaining a Driver's License in India: An Experimental Approach to Studying Corruption." *Quarterly Journal of Economics,* 122(4): 1639-76.

Burger, Jerry M. 2009. "Replicating Milgram: Would people still obey today?" *American Psychologist*, 64(1): 1-11.

Burtless, Gary. 1995. "The Case for Randomized Field Trials in Economic and Policy Research." *Journal of Economic Perspectives*, 9(2): 63-84.

Dearman, David T. and James E. Beard. 2009. "Ethical Issues in Accounting and Economics Experimental Research: Inducing Strategic Misrepresentation." *Ethics and Behavior*, 19(1): 51-9.

Elms, Alan C. 1972. *Social Psychology and Social Relevance*. Boston: Little, Brown.

Gugerty, Mary Kay and Michael Kremer. 2008. "Outside Funding and the Dynamics of Participation in Community Associations." *American Journal of Political Science*, 52(3): 585-602.

Harmon, Amy. 2010. "New Drugs Stir Debate on Rules of Clinical Trials." *New York Times*, September 19, 2010. http://www.nytimes.com/2010/09/19/health/research/19trial.html? _r=1&emc=eta1.

Harrison, Glenn W. and John A. List. 2004. "Field Experiments." *Journal of Economic Literature*, 42(4): 1009-55.

Heckman, James J. 1998. "Detecting Discrimination." *Journal of Economic Perspectives*, 12(2): 101-16.

Hegtvedt, Karen A. 2007. "Ethics and Experiments." in *Laboratory Experiments in the Social Sciences*, edited by Murray Webster and Jane Sell. Amsterdam: Elsevier.

Hertwig, Ralph and Andreas Ortmann. 2001. "Experimental Practices in Economics: A Methodological Challenge for Psychologists?" *Behavioral and Brain Sciences*, 24(3): 383-403.

Johnson, Eric J. and Daniel Goldstein. 2003. "Do Defaults Save Lives?" *Science*, 302(5649): 1338-9.

MacCoun, Robert J. and Norbert L. Kerr. 1987. "Suspicion in the Psychological Laboratory: Kelman's Prophecy Revisited." *American Psychologist*, 42: 199.

Madrian, Brigitte and Dennis Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics,* 116: 1149-1187.

Milgram, Stanley. 1963. "Behavioral Study of Obedience." *Journal of Abnormal and Social Psychology*, 67: 371–8.

Milgram, Stanley. 1965. "Some Conditions of Obedience and Disobedience to Authority." *Human Relations*, 18: 57–76.

Milgram, Stanley. 1974. *Obedience to Authority: An Experimental View*. New York: Harper & Row.

Nofsinger, John. 2009. "Opt-in and Opt-out Pension Design." Psychology Today, August 30, 2009. http://www.psychologytoday.com/blog/mind-my-money/200908/opt-in-and-opt-out-pension-design

Thaler, Richard. 2009. "Opting in vs. Opting Out." *New York Times*, September 26, 2009. http://www.nytimes.com/2009/09/27/business/economy/27view.html