

NBER WORKING PAPER SERIES

WELFARE IMPLICATIONS OF LEARNING THROUGH SOLICITATION VERSUS
DIVERSIFICATION IN HEALTH CARE

Anirban Basu

Working Paper 20376
<http://www.nber.org/papers/w20376>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2014

I am grateful for comments from Karl Claxton and David Meltzer and support from NIH research grants RC4CA155809 and R01CA155329. The opinions expressed here are those of the author only. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Anirban Basu. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Welfare Implications of Learning Through Solicitation versus Diversification in Health Care
Anirban Basu
NBER Working Paper No. 20376
August 2014
JEL No. C01,C9,D6,I1,I18

ABSTRACT

This paper uses Roy's model of sorting behavior to study welfare implication of current health care data production infrastructure that relies on solicitation of research subjects. We show that due to severe adverse-selection issues, directionality of bias cannot be established and welfare may decrease due to new data. Direct diversification of treatment receipt may solve these issues but is infeasible. Unifying Manski's work diversified treatment choice under ambiguity and Heckman's work on estimating heterogeneous treatment effects, the paper proposes a new infrastructure based on temporary diversification of access that resolves the prior issues and can identify nuanced effect heterogeneity.

Anirban Basu
Departments of Health Services,
Pharmacy and Economics
University of Washington
1959 NE Pacific St
Box - 357660
Seattle WA 98195
and NBER
basua@uw.edu

One of the fundamental challenges in health care markets is lack of information about the quality of medical care and technology (Arrow, 1963). Unlike other product markets, such as those of electronics and consumer appliances, when new health technologies are introduced in the market, minimal information about their comparative quality with respect to older technology or products is available to consumers, a.k.a. patients and physicians. It is often difficult to predict or anticipate such comparative quality information based on product attributes such mechanism of drug action, bioavailability or other mechanistic features because of the uncertainty in predicting the interaction of the mechanical, pharmacological or procedural features of a new medical technology with human biology.

Information on medical product quality is usually generated by employing an artificial form of '*learning by doing*' mechanism where a selected group of individuals (doers) are allowed to consume alternative medical products (e.g. using standard statistical designs, such as randomized assignment of patients to products). Wisdom from their experiences is disseminated to other individuals who will face the choice of using these medical products in the near future and also to inform social policies on access.¹ Most public and private stakeholders that are engaged in data production on medical quality signals have employed such mechanisms. Recently, large boluses of public investments were made in the US, under the umbrella term "comparative effectiveness research" (CER) and patient-centered outcomes research (PCOR),² to facilitate production of such data on alternative

¹ There are situations where learning from own's doing is popular, aka the repeated use of pharmaceutical products in chronic illnesses.

² Patient Protection and Affordable Care Act of 2009, H.R. 3590, 111th Congress §6301 (2010).

medical technologies that are currently being used in clinical practice, albeit with incomplete knowledge of their comparative qualities.³

In this paper, using a simple Roy's model (1951) of sorting behavior we prove that when medical treatments are available to patients under insurance, any data production infrastructure for comparative medical quality that relies on voluntary participation of subjects (which, unfortunately is and has been the norm for CER studies) fails to identify any interpretable treatment effect parameters and therefore fails to inform either the individual patient on *optimal* medical care use or a social insurer on *optimal* medical care insurance coverage.⁴

The implications of this finding is substantial. Comparative quality information influences individual's *ex-ante* perception of net benefits of treatment and forms the basis for the market demand curve (Basu 2011). Accurate estimates of this information carry tremendous welfare value as it reduces decision uncertainty for individual consumers and help them choose a product that will maximize recovery from an illness. Consequently, it helps align the market demand curve of a product with its normative demand curve that is based on true marginal benefits. In contrast, when comparative quality information generated by CER-type research is incomplete, it has the potential to misguide treatment choices since *ex-ante* perception of benefits do not coincide with the *ex-post* accrual of the same, and result in welfare losses. For example, inefficiencies in the choice of medical

³ Throughout our paper, we assume the CER compares two medical technologies that have been approved for use based on meeting the minimum safety thresholds as those set by the Food and Drug Administration of the United States. Our discussions do not encompass evaluation of experimental therapies. Such discussions are delegated to future work. Also see Philipson (1995).

⁴ Note that our assertions about optimality are very general and does not depend on specific welfare functions. What we prove is that the structural target parameters on which information is required to maximize any welfare function is not informed by current data production infrastructure.

products due to incomplete information can accentuate the inefficiencies due to moral hazard stemming from health insurance that separates demand prices from supply prices of medical products (Arrow, 1963; Pauly, 1968). Such added inefficiencies can translate to higher premiums and less protection against risk in both competitive and non-competitive insurance markets.

We begin in the next section by laying out the role of perfect and complete information on decision making and on outputs from the health care markets. In our paper, a social planner is not charged with prescribing treatment to each individual but is rather asked to decide on whether to pay for a treatment through insurance. Individuals, with their idiosyncratic and, perhaps, evidence-based, beliefs about treatment outcomes choose treatments. In Sections 2 and 3 we highlight the current data production infrastructure and prove why it would produce incomplete information. We study the implications for such incompleteness on decision making and welfare.

In Section 4, we introduce a new framework for data production that can efficiently resolve the biases inherent in the current data production infrastructure by using diversification of access to create a conduit for learning about meaningful and decision-relevant effect parameters. This work unifies two broad themes in the econometric literature, one based on Manski's work on treatment choice under ambiguity (Manski 2000, 2004, 2009) that utilizes the concept of diversification of treatment as posited in Manski (2009) and the other based on Heckman, Vytlačil and others' works on estimating heterogeneous treatment effects (Heckman 1997, 2001; Heckman and Vytlačil 1999, 2001; Heckman et al. 2006). We prove how this framework can help overcome inefficiencies in health care markets that stem from incomplete information. We discuss many implications of instituting such a data production framework such as how it can fix incentives for investment in data

production for public and private stakeholders in the long-run and its potential to fuel global completion in data production for medical quality.

I. The Role of Perfect Quality Information

Let us begin with a problem of evaluating the comparative effectiveness of a new (approved) treatment compared to a control/standard treatment for a population of N patients indexed by i . Standard treatment may also include the do-nothing option. Let the individual-level true treatment effect represent the benefits (net of harms) of the new treatment over the control and is denoted by b_i . Let p denote the price of the new treatment which is also the marginal cost for manufacturing the new treatment.⁵

Patients are members of risk classes Ω , $\Omega=1,2,..k$; $k \leq N$, which determine heterogeneity in treatment effects across individuals through a production function $b()$, i.e. $b_i = b(\Omega)$. Risk classes may be defined by the combination of various risk factors such as severity of illness, patient demographics, their genetic makeup, preferences and many others factors. In order to predict an individual-level treatment effect, one requires knowledge about both the patient's risk class Ω and the production function $b()$. CER is responsible for estimating the production function. Consider, for example without loss of generality, a production function that is expressed as a formulation of splines:

$$b_i = \sum_k \alpha_k \cdot I(\Omega_i = k) \tag{Eq. 1}$$

where $I()$ is an indicator function and α_k is interpreted as the comparative effect of the new treatment over the standard treatment in risk class k . Let's assume that this comparative effect is expressed in monetary terms. That is the effectiveness unit is

⁵ Assume for now that the marginal cost is constant.

multiplied with the some predefined threshold representing the monetary value of the marginal unit of benefit.⁶

A population level average effect parameter is given as

$$\Delta = \sum_k \Pr(\Omega = k) \cdot \alpha_k \quad \text{Eq. 2}$$

There are two types of decision makers, 1) the patient-physician dyad, which we will refer to as the individual decision maker, is assumed to always have knowledge about their risk class; and 2) an insurer or social planner who decides the coinsurance rate for providing health insurance coverage for the new treatment.

A. First-Best Scenario

Under complete information, both the insurer and the individuals are aware of the risk classes and the production function and are able to perfectly predict b_i . Here, individuals will choose treatment only if $b_i - p \geq 0$, when they are exposed to the full price of treatments. This is efficient, assuming that the generic social welfare function we will aim to maximize consumer surplus net of expenditure (i.e. a Net-Benefit criterion). If individuals had full insurance they would choose treatment if $b_i \geq 0$. Since the insurer can fully anticipate this individual behavior, she can provide full coverage for treatment only for those individuals who would experience benefits greater than cost and not provide coverage for the rest. Thus, there is no efficiency loss due to moral hazard. The question on economic evaluation about providing coverage on the new treatment can be answered in an individualized fashion based on individualized comparative effectiveness information.

⁶ Under the welfare economic foundations, this threshold is the inverse marginal utility of income (Weinstein and Zeckhauser, 1977; Garber and Phelps, 1997; Meltzer, 1997).

B. Second-Best Scenario

The traditional theory of health insurance recognizes that such a complete information scenario is not realistic. Specifically, it assumes that there exist asymmetry of information where, even though, individuals are assumed to be aware of Ω_k and $b()$ and to be able to combine them to predict b_i perfectly, the insurer cannot as they have either no or only partial information on Ω_k (Arrow 1963; Pauly and Blavin 2008). Consequently, the insurer cannot exclude patients from coverage who would get treatment benefits lower than the cost of treatment (i.e. $b_i - p < 0$). This leads to moral hazard (Pauly, 2008) and the insurer may offer coverage with a fractional coinsurance rate (r), which is the fraction of price a patient must pay in order to receive treatment. When $r = 1$, the new medical product is not covered through insurance.

We assume individuals choose treatment by maximizing a generic Net-Benefit criterion that is based on their perceived benefits from treatment net of the demand price they face in acquiring the treatment. We also assume that the social insurer's goal is to maximize consumer surplus as is realized *ex post* based on individual level choices. Therefore, throughout this paper, we will express the realized population level benefits under different levels of coverage for the new treatment as changes to the total outcomes had all patients taken the standard treatment. Under any co-insurance rate r , $r \in [0, 1]$, this population level benefits, H^0 , is given as

$$H^0 = \sum_i \sum_k I(\alpha_k - r \cdot p \geq 0) \cdot \alpha_k \cdot I(\Omega_i = k). \quad \text{Eq. 3}$$

That is, when individuals have complete information they choose to receive the new treatment only if $\alpha_k \geq r \cdot p$. The population level benefit is then expressed as an aggregation of α_k across those individuals. Note that under the second-best

scenario, individuals who would expect to get harmed by treatment (i.e. $\alpha_k < 0$) would not select treatment even if it were available to them for free, thereby self-limiting the magnitude of moral hazard.

For a social insurer's point of view, an optimal co-insurance rate may be expressed as a solution to maximizing H^0 net of costs and taking into account the social value of risk protection provided by the insurance, $V_{2nd}(r)$:

$$r^* = \arg \max_r (V_{2nd}(r) + \sum_i \sum_k I(\alpha_k - r \cdot p \geq 0) \cdot (\alpha_k - p) \cdot I(\Omega_i = k)) \quad \text{Eq. 4}$$

In equilibrium, the welfare loss due to moral hazard due should equate to the social value of risk protection (Manning and Marquis, 1996). Consequently, the moral hazard (welfare loss) under optimal coinsurance rate in a second-best scenario is given as

$$L_{2nd}(r^*) = \left(\sum_i \sum_k I(r^* \cdot p \leq \alpha_k < p) \cdot (p - \alpha_k) \cdot I(\Omega_i = k) \right) = \sum_{\alpha_k=r^* \cdot p}^p N_k \cdot (p - \alpha_k), \quad \text{Eq. 5}$$

which constitutes the welfare loss due to the total number of individuals in each risk group (N_k) who would choose treatment given the lower demand price ($r \cdot p$) but ultimately obtain benefits lesser than the price of treatment, i.e. $r^* \cdot p \leq \alpha_k < p$.

II. Data Production and Incompleteness in Quality Information

Reality, however, deviates from both the first and second best scenarios, because both individual and the social decision makers face incomplete comparative information. To understand this incompleteness, one must study the data production mechanisms in place. We consider and compare the circumstances before and after a CER study. We begin by understanding the consequences of

incomplete information before a CER is conducted and why added investments for data productions, such as those provisions by the latest legislations, are called for. We then study how the current mechanisms of CER may continue to propagate and even enhance the welfare losses due to incomplete information.

A. Pre-CER information and choices

Before CER is conducted it is safe to assume that α_k is not known with certainty both at the individual and the societal level. However, prior knowledge, obtained from evidence (of size n) generated during the process of approving the use of this new medical product would determine an individual patient's anticipated belief about the incremental benefits of treatment given one's own risk class. Let this evidence suggest that the average effect of treatment is $\bar{\alpha}$ that is a random draw from $\text{Normal}(\Delta, \sigma_\Delta^2/n)$,⁷ where Δ is the average effect parameter defined in (2) and σ_Δ is the heterogeneity and the standard deviation of the effect in the population. Let individual beliefs, α'_i , be given as a single draws from the distribution $\text{Normal}(\bar{\alpha}, s^2)$ where s is the estimated standard deviation from prior evidence. It is assumed that s^2 is a consistent estimator of σ_Δ^2 . The schedule of α'_i across individual patients determines the marginal benefits curve in the population in the absence of a CER. Moreover, the social insurer may not have perfect information on either α_k or α'_i . However, she may have information about the average effect, $\bar{\alpha}$. The best a social insurer can do at this point is to calculate the average net monetary benefits of treatment,

⁷ We take a conservative approach is assuming that $\bar{\alpha}$ and α'_i are consistent estimators of Δ . To the extent this is not true, the welfare losses described below may be higher.

$$\bar{\alpha} - p$$

Eq. 6

and recommend coverage if $\bar{\alpha} - p \geq 0$.⁸

Without loss of generality,

Assumption 1: Let $\Delta > 0$, the true population average treatment effect is positive, but α_k 's span the whole real line.

Assumption 2: Let $\bar{\alpha} - p \geq 0$ and full coverage was recommended, i.e. $r^* = 0$.

Theorem 1: Under Assumptions 1 and 2, $L_{PRE}(0) > L_{2nd}(r^*)$ for $\forall r^* \in [0, 1]$ if $\alpha'_i \geq 0, \forall i$. The welfare loss under pre-CER information with full insurance coverage is strictly larger than the welfare loss under any second-best scenario as long as all individuals perceive a positive benefit from treatment.

Proof: Under the Pre-CER scenario, the welfare loss is due to two groups of people making inefficient choices. The first group consists of people who fail to receive treatment because their $\alpha'_i < 0$ but they belong to risk group where the treatment produces incremental benefits that are more than the price of the treatment (i.e. $\alpha_k > p$). The second group consists of individuals who would consume the medical product but obtain a benefit less than its price. Therefore total welfare loss is given by:

$$L_{PRE,CER} = \left(\sum_i \sum_k I(\alpha'_i < 0) \cdot I(\alpha_k - p > 0) \cdot (\alpha_k - p) \cdot I(\Omega_i = k) \right) \\ + \left(\sum_i \sum_k I(\alpha'_i \geq 0) \cdot I(\alpha_k - p < 0) \cdot (p - \alpha_k) \cdot I(\Omega_i = k) \right)$$

⁸ This is, in fact, the standard method used in most cost-effectiveness modeling studies that try to evaluate the cost-effectiveness of a new approved treatment for which there is no head-to-head comparison with its alternatives.

$$= \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k=-\infty}^p \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (p - \alpha_k), \quad \text{Eq. 7}$$

where $\Phi(\cdot)$ is a cumulative normal distribution. Note that the second expression in (7) would include all individuals who would obtain negative benefit from consuming the medical product (i.e. $\alpha_k < 0$) but are led to believe that they would get a positive benefit ($\alpha'_i \geq 0$) due to incompleteness in information. Comparing (7) to (5),

$$\begin{aligned} L_{PRE,CER}(0) - L_{2nd}(0) &= \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (\alpha_k - p) \\ &+ \sum_{\alpha_k=0}^p \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) - 1 \right\} \cdot N_k \cdot (p - \alpha_k) \\ &+ \sum_{\alpha_k=-\infty}^0 \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (p - \alpha_k) \end{aligned} \quad \text{Eq. 8}$$

The first and the third terms in (8) are the incremental losses due to incomplete information pre-CER. The first term is the same as in (7) and comprises of individuals who fail to take treatment but would benefit more than its price. The loss represented in the third term emanate from risk groups where $\alpha_k < 0$ and a fraction of individuals in these risk groups take treatment based on their perceived benefits, which was not the case under the second-best scenario.

The second term in (8) is a pervasive *benefit* of incomplete information compared to the second-best scenario (expressed as negative loss). The benefit emanate from risk groups where $0 < \alpha_k < p$ and a fraction of individuals in these risk groups forgo treatment based on their perceived benefits (i.e. their $\alpha'_i < 0$), which was not the case under the second-best scenario, thereby generating welfare gains.

Under assumptions 1 and 2, if one assumes that $\alpha'_i \geq 0, \forall i$, that is every individual perceives a positive benefit from treatment, Theorem 1 is proved from (8) as the second term drops out and $L_{PRE}(0) > L_{2nd}(0)$. Thus, naturally, $V_{PRE}(0) > V_{2nd}(r^*)$ for $\forall r^* \in [0, 1]$. ■

B. An ideal role for CER

Often an ideal CER is construed as one having larger sample size. In fact much of the value of information literature in medicine has focused estimating the marginal value of a trial with additional patients enrolled (see literature on the Expected Value of Sample Information, EVSI). However, it is not clear whether, in the presence of heterogeneity, such an approach to CER, is welfare enhancing. For example, as $n \rightarrow \infty$, $\bar{\alpha} \xrightarrow{p} \Delta$, $s^2 \xrightarrow{p} \sigma_\Delta^2$. Consequently, the welfare loss post CER of infinite sample will be:

$$L_{POST,CER(n \rightarrow \infty)} = \sum_{\alpha_k = p}^{\infty} \left\{ 1 - \Phi\left(\frac{\Delta}{\sigma_\Delta}\right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k = -\infty}^p \left\{ \Phi\left(\frac{\Delta}{\sigma_\Delta}\right) \right\} \cdot N_k \cdot (p - \alpha_k),$$

Eq. 9

Note that if $\Phi\left(\frac{\Delta}{\sigma_\Delta}\right) < \Phi\left(\frac{\bar{\alpha}}{s^2}\right)$, a CER of infinite size will decrease the magnitude of loss in the second term but increase the magnitude of loss in the first term as compared to L_{PRE} and vice versa. Therefore, the value of such a CER study is indeterminate.

An ideal role for CER would be when, even though the social insurer continue to believe in the average effect $\bar{\alpha}$, it can enable individual belief, α'_i to be a draw from the distribution $\text{Normal}(\alpha_k, \sigma_k^2)$ where σ_k^2 is variance for risk-group-specific

effects.⁹ If coverage is recommended based on the criteria in (6), welfare losses under post ideal-CER scenario will be a modification of (7) to

$$L_{POST^*} = \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi\left(\frac{\alpha_k}{\sigma_k}\right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k=-\infty}^p \left\{ \Phi\left(\frac{\alpha_k}{\sigma_k}\right) \right\} \cdot N_k \cdot (p - \alpha_k) \quad \text{Eq. 10}$$

Since, $\Phi\left(\frac{\alpha_k}{\sigma_k}\right) > \Phi\left(\frac{\bar{\alpha}}{s}\right)$ for $\alpha_k > p$ as by construction $\alpha_k > \bar{\alpha}$ for all $\alpha_k > p$ and

$\Phi\left(\frac{\alpha_k}{\sigma_k}\right) < \Phi\left(\frac{\bar{\alpha}}{s}\right)$ for $\alpha_k \leq p$ as by construction $\alpha_k < \bar{\alpha}$ for all $\alpha_k \leq p$,¹⁰

$L_{POST^*}(r) < L_{PRE}(r) \quad \forall r^* \in [0, 1]$. This unambiguous dominance of an ideal CER over pre-CER scenario arises because individuals are able to better self-select their optimal treatment based on the risk group specific knowledge generated from an ideal CER. In fact, as $\sigma_k \rightarrow 0$, $\Phi\left(\frac{\alpha_k}{\sigma_k}\right) \rightarrow 1$ for $\alpha_k \geq 0$ and $\Phi\left(\frac{\alpha_k}{\sigma_k}\right) \rightarrow 0$ for $\alpha_k < 0$.

Consequently, $L_{POST^*}(r) \rightarrow L_{2nd}(r)$. Therefore, one can potentially approach a second-best scenario under any level of insurance coverage if new CER studies are able to generate information that can enable individuals to better self-select treatments based on their risk classes, even if the social insurer is unaware of these heterogeneous effects. The growing awareness of the potential value of such has led to considerable federal investment in CER. New legislation has also identified the need to risk stratify comparative effectiveness.

However, the current data production infrastructure for CER may not be aligned with the goals of such legislation. The gold standard of data production in medical care involves controlled experiments, where alternative treatments under investigation are allocated to a selected group of patients by a chance mechanism.

⁹ In practice, even in the absence of CER such a situation may arise, when individuals learn by repeated consumption of therapy (e.g. pharmaceuticals) or physicians are able to anticipate effect heterogeneity based on baseline risks.

¹⁰ Assuming $\sigma_k = s$ for all k .

We will refer to such a mechanism as a randomized clinical trial (RCT) henceforth. We consider two issues within this data generating infrastructure that contributes towards the inability of current CER infrastructure to resolve incompleteness in information: selection in RCT enrollment and target parameters for RCTs.

C. Non-ideal design and implementation of CER studies – understanding selection into randomized trials

Unlike evaluation of experimental therapy where enrollments may be more likely driven by altruistic motives, CER and economic evaluation is about approved and existing therapies available to patients. Therefore, there must be a strong implicit selection process for patients who provide consent to enroll in a CER RCT, in which they have a non-trivial probabilistic expectation of receiving a treatment (most likely the new treatment) that they have some difficulty in accessing outside of the RCT. Such difficulties must be because the cost of accessing the new treatments outside RCT must be high either due to differential insurance coverage of the treatments or due to strong physician preferences for one therapy over other. Consequently, this selection process implies that the anticipated expected incremental benefits for patients who enroll in a CER RCT must be *less* than the cost of acquiring the treatment outside of RCT.¹¹ Otherwise they would have taken the treatment without enrolling in an RCT. In other words, patients who anticipate large incremental benefits or incremental harms from new treatment compared to standard care are less likely to enroll in RCTs.

This brings to question whether the anticipated benefits of treatment (α'_i) are related in any form to the true benefits of treatment b_i , even in the absence of formal

¹¹ Often enrollees are paid a monetary sum to compensate for their time spent participating in the RCTs.

CER. Obviously, if $\alpha'_i = b_i$, then the value of any additional CER becomes zero as each patient already know their true benefits. On the other hand, the value of CER is maximized when $\alpha'_i \perp b_i$, where \perp denotes statistical independence. In practice, however, it is not uncommon to find some dependency between α'_i and b_i . Such dependencies may arise from biological knowledge about the treatment's mechanism of actions, past experiences by physicians on using similar treatments on certain patient risk-groups and by patient's own learning by doing mechanism in a chronic disease setting. Under such dependencies, effect of selection into RCT becomes non-trivial. Specifically, we show

Theorem 2: *A CER randomized trial produces an unbiased estimate of the population average treatment effect Δ (in Eq. 2) if and only if $\alpha'_i \perp b_i$. If $\text{Corr}(\alpha'_i, b_i) > 0$, RCTs will typically find small positive benefits of treatment.*

Proof: We formalize selection into a CER RCT following Roy's model (1951) of self-selection using the following notation

$$S_i = I(U_i^* \geq 0), \tag{Eq. 11}$$

where S is an indicator for enrolling in an RCT that is driven by the latent utility U^* for enrolling. Again, without loss of generality, U_i^* is interpreted as the anticipated incremental net benefits (net of costs) of enrolling versus not enrolling in an RCT for individual i given that the individual anticipates a positive benefit from treatment (i.e. $\alpha'_i > 0$)¹²

$$U_i^* = (\pi_R \cdot \alpha'_i - C_{RCT}) - (\alpha'_i - C_{OUT}) = (C_{OUT} - C_{RCT}) - (1 - \pi_R) \cdot \alpha'_i, \text{ if } \alpha'_i > 0 \tag{Eq. 12}$$

¹² If individual anticipates a negative benefit from treatment he would not consider enrolling in the first place.

where C_{RCT} and C_{OUT} are the costs of accessing the treatment within and outside an RCT respectively; π_R is the known random probability of receiving the new treatment within the CER RCT.

In the presence of uncertainty, the population probability of an RCT enrollment is given by

$$\pi = \Pr(S_i) = E(U_i^* \geq 0) = \Pr(0 < \alpha'_i < (C_{OUT} - C_{RCT}) / (1 - \pi_R)) \quad \text{Eq. 13}$$

Therefore, only patients with anticipated benefits are positive and less than the expected incremental cost of accessing treatment outside RCT would enroll. Interestingly, when $C_{OUT} \approx C_{RCT}$, enrollment in CER can be quite difficult. Similarly, as π_R decreases, it reduces the cost differential between accessing the new treatment outside and within the RCT, thereby lowering the probability of RCT enrollment. These factors severely limit the generalizability of results from CER RCTs. For example, in one of the few surveys ever conducted to understand the factors that determine RCT enrollment, it was found that only 2.7% of eligible patients enrolled in clinical oncology trials (Movsas et al. 2007).

Target Parameters for RCT The target for most RCTs, if not all, is to estimate an average effect of treatment compared to the control among the RCT enrollees. However, keeping aside the challenges of implementing a protocol driven agenda, such an average effect is a weighted average of risk-class-specific effects where the weights are arbitrarily defined based on the risk class-specific propensity to enroll in the RCT. Therefore, the target parameter for RCT is given by:

$$\Delta_{RCT} = \sum_k w_k \cdot \alpha_k \quad \text{Eq. 14}$$

where the weights $w_k = (\pi_k \cdot \Pr(\Omega = k)) / \sum_k (\pi_k \cdot \Pr(\Omega = k))$ and

$\pi_k = \Pr(0 < \alpha'_i < (C_{OUT} - C_{RCT}) / (1 - \pi_R) \mid \Omega_i = k)$. The degree of selection in the trial determines the target parameter of an RCT.

When $\alpha'_i \perp b_i$, $F(\alpha'_i | \Omega_i = k) = F(\alpha'_i) \forall k$, where $F()$ is the cumulative distribution function. This implies, $\pi_k = \pi, \forall k \Rightarrow w_k = \Pr(\Omega = k), \forall k \Rightarrow \Delta_{RCT} = \Delta$ (according to Eq (2)). On the contrary, $\Delta_{RCT} \neq \Delta$ if $\alpha'_i \not\perp b_i$, since the weights would vary depending on which risk classes are more likely to enroll in the RCT.

In fact, $\underline{Corr(\alpha'_i, b_i)} > 0 \Rightarrow Corr(w_k, \alpha_k) < 0$ for $\alpha'_i > 0$ and $Corr(w_k, \alpha_k) > 0$ for $\alpha'_i \leq 0$. Individuals who correctly anticipate large positive or negative benefits from treatment are less likely to enroll in RCTs. In fact, Eq (13) suggest that the margin of individual who enroll in RCT anticipates a moderated positive magnitude of benefits from treatment. Given that $Corr(\alpha'_i, b_i) \geq 0$, it implies that RCT results would typically find small positive benefits of a newer treatment and the generalizability of these results to the whole target population remains severely compromised.¹³ ■

Consequently, in the presence of any anticipatory knowledge about true treatment effects, the average effect from an RCT is not a consistent estimator for either population average effect or the average effect of any segment of the population: $E(\hat{\Delta}_{RCT}) \neq \Delta$ and $E(\hat{\Delta}_{RCT}) \neq \alpha_k \forall k$. Next, we study how such results can mislead individual level decision making and create inefficiencies both through population-level coverage decisions and individual treatment selections.

Assumption 3: In what follows, we will assume $Corr(\alpha'_i, b_i) > 0$ even in the absence of a formal CER.

¹³ It is possible that under an ideal symmetric condition, the weights are such that equivalent portions of the risk-groups with large positive effects and those with large negative effects select out of enrolling and the average effect among the enrollees still reflects the population average. However, such a scenario is highly unlikely.

D. Implications of incompleteness for decision making

Theorem 3: (a) Under Assumption 3, CER RCT may misguide a social planner to provide coverage on treatments with negative average net health benefits and to withhold coverage on treatments with positive average net health benefits.

(b) Under Assumptions 3, $L_{POST}(0) > = < L_{PRE}(0)$. The welfare loss under post-CER information can be larger than that under pre-CER scenario with full insurance coverage.

Proof: (a) Based on CER RCT results, the social planner updates her belief over the average effect of the new treatment using a Bayesian updating rule (Basu et al. 2011):

$$\bar{\alpha} = \theta \cdot \bar{\alpha} + (1 - \theta) \cdot \hat{\Delta}_{RCT} \tag{Eq. 15}$$

where the weight θ is determined by a weighted average of prior uncertainty σ_{Δ}^2 and the sampling variance of $\hat{\Delta}_{RCT}$, and calculates the average net monetary benefits of treatment to be $\bar{\alpha} - p$. Under Assumption 3, Theorem 2 proves that $E(\hat{\Delta}_{RCT}) > 0$ but $E(\hat{\Delta}_{RCT}) > = < \Delta$. Therefore, since $E(\bar{\alpha}) = \Delta$, $E(\bar{\alpha} - p) < E((\Delta - p))$ if $E(\hat{\Delta}_{RCT}) < \Delta$ and $E(\bar{\alpha} - p) > E((\Delta - p))$ if $E(\hat{\Delta}_{RCT}) > \Delta$.

This implies that CER RCT may misguide a social planner to provide coverage on treatments with negative average net health benefits and to withhold coverage on treatments with positive average net health benefits. This also highlights the fact that economic evaluations based on CER studies can be misleading. ■

(b) Individual beliefs, α_i'' , about comparative effects following the CER RCT using a similar Bayesian updating rule (Basu et al. 2011):

$$\alpha_i'' = \theta_i \cdot \alpha_i' + (1 - \theta_i) \cdot \hat{\Delta}_{RCT} \tag{Eq. 16}$$

where the weights θ_i are determined by a weighted average of prior uncertainty s^2 and the sampling variance of $\hat{\Delta}_{RCT}$. It is important to note that even though original beliefs may have been consistent, i.e., $E(\alpha'_i) = b_i$, after CER, $E(\alpha''_i) \neq b_i$. Most importantly, $\alpha''_i > \alpha'_i$ if $\alpha'_i < 0$, under Assumption 3, since $E(\hat{\Delta}_{RCT}) > 0$. That is, some patients who would have had originally anticipated a negative effect from treatment, may be rightfully so, are now led to believe in a larger, presumably, positive effect from treatment. Similarly, patients who would have, rightfully anticipated large benefits from treatment, would have their updated anticipation moderated by the small effect size estimated in the RCTs. Thus the average result from a CER study that is based on voluntary participation actually misleads individuals about their own comparative effectiveness. The welfare loss with the post CER information is given by:

$$L_{POST,CER} = \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi\left(\frac{\mu_k(\alpha''_i)}{s_k(\alpha''_i)}\right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k=-\infty}^p \left\{ \Phi\left(\frac{\mu_k(\alpha''_i)}{s_k(\alpha''_i)}\right) \right\} \cdot N_k \cdot (p - \alpha_k),$$

Eq. 17

where $\mu_k(\alpha''_i) = E(\alpha''_i | \Omega_i = k)$ and $s_k^2(\alpha''_i) = Var(\alpha''_i | \Omega_i = k)$. Since $\frac{\mu_k(\alpha''_i)}{s_k(\alpha''_i)} < = > \frac{\bar{\alpha}}{s}$

for any k , it proves that welfare loss under post-CER information can be larger than that under pre-CER scenario with full insurance coverage. Only when $\frac{\mu_k(\alpha''_i)}{s_k(\alpha''_i)} \geq \frac{\bar{\alpha}}{s}$

for $\alpha_k \geq p$ and $\frac{\mu_k(\alpha''_i)}{s_k(\alpha''_i)} < \frac{\bar{\alpha}}{s}$ for $\alpha_k < p$, the CER infrastructure is welfare enhancing

($L_{POST,CER} < L_{PRE,CER}$).

III. Learning through Diversification (LtD): A New Framework for Data Production

As we have shown in the previous sections, current CER framework that relies on voluntary participation fails to consistently inform either the population-level or individual-level comparative effect parameters and cannot potentially lead us towards the second-best solutions (in fact, it may lower welfare through evidence-based misguidance).

Manski (2009) proposed that one way a social decision maker can maximize welfare is through fractional allocations, where a random fraction of the patient population received one treatment while the other receives the alternative. Manski argues that, given the ambiguity of evidence on counterfactual outcomes, such an allocation would maximize a broad set of utilitarian welfare function for the social decision maker. Manski (2009) also points out that such an allocation automatically creates randomized experiments, which are particularly important for learning treatment responses. The current proposal builds on this idea of “diversified treatment” proposed by Manski (2009). However, our proposal takes into account two realities in the context of health care.

First is that it is almost impossible, at least in the United States, to completely restrict “receipt” of a treatment that has crossed the regulatory and evidentiary hurdles and has been approved on the basis of safety and efficacy. Therefore diversification of treatment allocation in terms of “receipt”, which is essential to answer CER and PCOR type question, is usually not possible.

Second, the social decision maker in the context of health care is typically involved on deciding on insurance coverage of medical treatment, while individual subjects are typically left to decide on the choice of treatment given insurance coverage. Therefore, a social decision maker’s problem can be viewed to be a two-step process (Deheja 2005). Under any information set, first physician decides

whether to prescribe treatment for each individual. Second, given this allocation, the social decision maker decides on the level of coverage for treatment. To the extent that one can combine the ideas of diversified treatment for the purpose of learning to that of the two-step process of social decision making on optimal coverage, one can improve the decision making for both the individual patients and the social decision maker. This is what the “Learning through Diversification” (LtD) infrastructure seeks to accomplish and potentially mimic the ideal CER designs discussed in Section II-B.

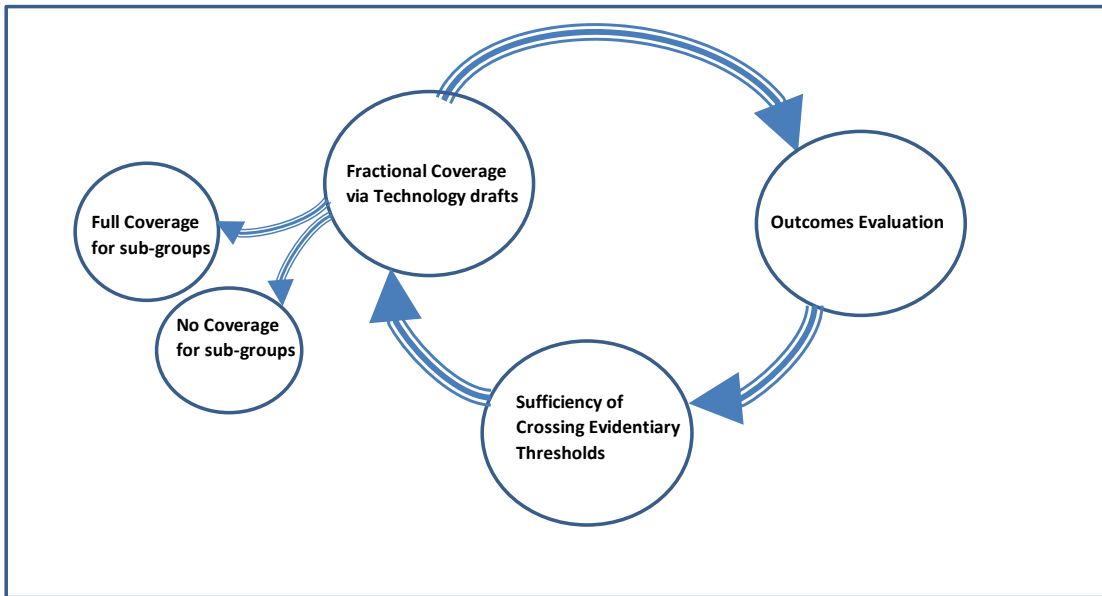


Figure 1: The Learning through diversification (LtD) infrastructure.

The idea behind this new data production infrastructure for CER is illustrated in Figure 1 and can be expressed in the following bullet points:

1. **Fractional Coverage** can be achieved using a technology lottery, \mathcal{D} : For *each new product*, develop a random order based on birth dates so that this new product with uncertain effectiveness profile will be paid at varying levels by insurance in the first year. That is, such coverage creates a

completely stochastic distribution of co-insurance rates in the population, $F(b_i | \mathbb{D}) = F(b_i)$. Note that the lottery is done anew for each new technology so that the probability that any one person would be denied coverage for all new technologies will approach zero with increasing number of technologies. Thus, the initial decision by a social insurer follows the idea proposed by Manski (2009) to make fractional treatment allocation. However, instead a binary assignment of access, we use a continuously varying cost of access, which will be important for assessing treatment effect heterogeneity (as shown below).

2. **Outcomes Evaluation:** Using the randomization inherent in the lottery, evaluating patient outcomes across different levels of coinsurance rates will directly answer the economic evaluation questions on expanding coverage for the target population. Additionally, the lottery would serve as a *perfect instrumental variable* to study the comparative effectiveness of receiving the new product versus its competitor and the heterogeneity in these effects in the population. To the extent such analyses can discover nuanced subgroups with higher than average benefits, subgroup-specific economic evaluations can be conducted to examine whether full coverage is warranted in these sub population. Obtaining precise information on comparative effectiveness within specific subgroups can itself drive demand to be selective in a positive way.
3. **Sequential Decision Making:** Based on the outcome evaluation results, fractional allocation rules can be adapted over time for specific risk groups. For example, among risk groups where estimated $\hat{\alpha}_k$ is expected to be positive above a certain evidentiary threshold may receive full payments for the treatment under the insurance and vice versa. Fractional allocation would continue within risk groups where ambiguity persists.

A. Key Features of the Learning through Diversification Infrastructure

Coinsurance (demand price) as an instrument

The outcomes evaluation part of the LtD framework employs the lottery as an instrumental variable (IV) to study comparative effectiveness of the new technology compared to controls and also explore treatment effect heterogeneity (Heckman 1996; Heckman 1999, 2001; Heckman and Vytlacil, 1999; Heckman et al., 2006, Basu 2012). A good instrumental variable must be orthogonal to the confounders (i.e. they are not contaminated) and is a strong predictor of treatment receipt. Traditional IV analyses focus around the debate on whether a chosen instrument is contaminated, given that the strength of the instrument is testable. In the LtD framework, the lottery, by design, is orthogonal to all confounders and therefore side-steps the typical debates in this literature. What is more interesting is that the strength of the instrument is driven by variation in out-of-pocket payments by patients that in turn will depend on the market price of the new technology and the price elasticity of demand. Although this is a testable feature of analyses, two issues are worth pointing out. First, the LtD infrastructure would be most efficient in data production for CER for new technologies that come at a high price tag, which aligns with the notion that most welfare can be generated if we can properly identify people would and would not benefit from the most expensive technologies. Second, the response of manufacturers of expensive technologies to artificially reduce price in order to undermine the LtD infrastructure may act against their bottom-line as they would send a wrong signal about the quality of their product and also delay production of evidence that would enable them to demand a higher price from payers.

The target parameters in the LtD infrastructure

The LtD infrastructure can be used to generate consistent estimates of both α_k and Δ by following the theory and methods of local instrumental variable approaches (Heckman and Vytlacil 1999, 2001; Heckman et al. 2006, Basu 2014).¹⁴ In summary, the probability of treatment choice given the lottery can be represented by $p(\mathcal{D}, \Omega)$ and estimated from data at hand. For any outcome, Y , under regular IV assumptions, the marginal treatment effects (Heckman and Vytlacil 1999) are identified by

$$\frac{\partial E_g(Y | \mathcal{D}, \Omega)}{\partial p} = E((Y_1 - Y_0) | \Omega, V = v) = MTE(\Omega, v), \quad \text{Eq. 18}$$

where $Y = D*Y_1 + (1 - D)*Y_0$ is the observed outcome and $V \sim \text{Uniform}[0,1]$ by construction.

MTE is perhaps the most nuanced estimable effect. It identifies an effect for an individual who is at the margin of choice such that one's levels of Ω and \mathcal{D} are just balanced by one's level of unobserved factors V , i.e. $P(\mathcal{D}, \Omega) = V$. Basu (2014) extends the LIV methods to identify Person-centered treatment (PeT) effects, which, for persons who choose treatment, follow

$$\begin{aligned} & E_{V|\Omega, P(\mathcal{D}), D} E(Y_1 - Y_0 | \Omega, P(\mathcal{D}), D = 1) \\ &= E(Y_1 - Y_0 | \Omega, V < P(\mathcal{D})) = P(\mathcal{D})^{-1} \int_0^{P(\mathcal{D})} MTE(\Omega, v) dv \end{aligned} \quad \text{Eq. 19}$$

Similarly, conditional effect for a person who did not choose treatment is obtained by integrating MTEs over values of V greater than p . Mean treatment

¹⁴ Note that it is important to pay close attention to dealing with essential heterogeneity within the LtD infrastructure. This is because treatment receipt will be correlated with factors such as income (because of the price differentials that the lottery creates), which in turn may be correlated with gains and losses from the new treatment.

effect parameters, α_k (Eq. 1) or Δ (Eq. 2) are readily obtained by average PeT effects over respective subgroups (Basu 2014). Thus the LtD structure can be used to recover consistent estimates of the decision relevant parameters.

Welfare effects

Let's take a two period model in which the first period is the Pre-CER period during which a CER study is being conducted. As the end of the first period the CER study results are disseminated and therefore the second period represents the post-CER world. Therefore, under assumptions 1, 2 and 3, total welfare loss over the two periods in a CER-based data production world is given as:

$$L_{CER} = L_{PRE,CER} + L_{POST,CER} \quad \text{Eq. 20}$$

where $L_{PRE,CER}$ and $L_{POST,CER}$ are given in Eq(7) and Eq(17) respectively.

Under the LtD framework of data production, welfare losses in both periods will be different. Let the total welfare loss over the two periods in a LtD-based data production world is given as:

$$L_{LtD} = L_{PRE,LtD} + L_{POST,LtD}$$

Since the LtD infrastructure allows for consistent estimation of the mean treatment effect parameters, in the second period, subjective beliefs about the benefits of treatment will align with the true values for subjects in each risk groups, $E(\hat{\alpha}_k) = \alpha_k$

. Moreover, given that these estimates were generated using data at the population scale, $\Phi\left(\frac{\hat{\alpha}_k}{\hat{\sigma}_k}\right) \rightarrow 1$ for $\alpha_k \geq 0$ and $\Phi\left(\frac{\hat{\alpha}_k}{\hat{\sigma}_k}\right) \rightarrow 0$ for $\alpha_k < 0$. Consequently,

$$L_{POST,LtD} = \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi\left(\frac{\hat{\alpha}_k}{\hat{\sigma}_k}\right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k=-\infty}^p \left\{ \Phi\left(\frac{\hat{\alpha}_k}{\hat{\sigma}_k}\right) \right\} \cdot N_k \cdot (p - \alpha_k) \rightarrow 0$$

Eq. 21

Therefore, an LtD infrastructure will always be welfare enhancing compared to the CER infrastructure as long as:

$$L_{LtD} - L_{CER} = (L_{PRE,LtD} + L_{POST,LtD}) - (L_{PRE,CER} + L_{POST,CER}) < 0$$

$$\rightarrow (L_{PRE,LtD}) < (L_{PRE,CER} + L_{POST,CER})$$

That is fractional allocation should be designed in a way that the loss during the initial data generation process is not greater than the combined losses under the CER infrastructure both during and after the data generation process. For the most power for analyses and consistency within the LtD infrastructure, it may be useful to set the mean coinsurance rate to be 0.5. The welfare losses, if any, during the data production period (that may can one or few years) can be easily recuperated from the welfare gains from LtD in the post data production period that typically lasts for many years.

Requirements of insurance system for implementation of the LtD infrastructure

It is natural for the LtD infrastructure to thrive within nationalized health care systems and integrated health care delivery systems to prevent attrition of patients. However, even in a competitive health insurance market an LtD infrastructure can be implemented when a central body helps generate the lotteries and all the private and public payers offer coverage following the lottery. In fact, no one payer would find it to be in their interest to break away from this commitment as otherwise it would invite adverse selection into their plan due to the providing generous coverage of new technologies.

A new era for outcomes research

The power of the LtD infrastructure not only allows the study of one CER question by a central body. It opens up a new era of outcomes research for all applied researchers. Any prospective or retrospective observational study evaluating a new technology in terms of either safety or effectiveness or both using novel patient-centered outcomes would benefit from the power of the lottery to generate variations that would enable them to make causal inferences (Heckman 1996). The LtD infrastructure can have the sufficient sample size to test heterogeneous treatment effects across nuanced subgroups, which is impossible within a traditional trial framework. Moreover, with the growth of people with multiple comorbid diseases, the LtD infrastructure can provide precise answers to how different treatments interact with each other and how to create the optimal bundle of treatment strategies in order to treat complex patients. Thus the LtD infrastructure has the potential to impact almost all fields of study in social sciences, public health and medicine, and help in the formulation of payment reforms that are based on episodes of care.

The expanded role of Health IT

Naturally, strong synergies exist between the LtD infrastructure and the Health IT systems that are designed to capture detailed data. Large national electronic databases are already being assembled in the context of CER. However, serious challenges remains as to how to detect signals from such databases based on which patient care can be changes. A fundamental problem is the identification of causal effects from such data. The LtD infrastructure can naturally lend tremendous identification properties for all types of evaluative questions.

IV. CONCLUSIONS

New medical treatment often gets approved based on its potential safety profile and its incremental efficacy compared to either placebo or a basic control treatment. Often superiority of the new treatment is not established and its comparative effectiveness compared to current clinical practice remains ambiguous. Nevertheless, the treatment is introduced for consumption at a substantial price in anticipation of a positive effectiveness claims based on efficacy results. Variability in effectiveness profile remains far from known. Under such ambiguity, a social insurer faces the challenge of deciding whether to pay for the treatment. In the US, public health insurance provider like the Medicare usually extend full coverage of these new treatments as long as there is positive efficacy signals. In UK and other countries, that formally look at the budget impact of coverage by comparing the costs of treatments (inclusive of its price) to the projected effectiveness based on efficacy signals. When coverage is allowed, a large welfare loss may ensue even when the new treatment can genuinely produce higher effectiveness in a certain margin of the population. This is due to the lack of evidence of how to match patients to alternative treatments.

In this paper, I show that under the status quo policy of extending coverage to a new treatment in the absence on complete information on its effectiveness profile, welfare loss can be substantial. These losses can be minimized by investments in studies that aims at generating such evidence. However, I also show, following a Roy's model of sorting behavior, that the current infrastructure on data production for this purpose, suffer for severe self-selection issues since the incentives to enroll in research studies is eroded by the low demand prices of obtaining medical care outside these studies. Consequently, the parameters identified from these studies do not inform any of the decision-relevant

parameters, either at the individual or the population level. I show that if one takes the normative approach of social insurer who is forward looking and wants to maximize any given social welfare function over a duration of time period (typically over the longevity of the new technology being considered), then it makes sense for the social insurer, irrespective of what coverage decision is made today, to device ways to learn about variations in incremental effectiveness of treatment in the population so that she can encourage/discourage appropriate subgroups to uptake/discard the new treatment. In fact, generation of such public evidence can directly inform individuals within the population to use this new treatment appropriately without additional effort by the social insurer, thereby approaching the second-best solutions. Based on this normative framework, I propose a positive Learning through Diversification (LtD) infrastructure, through which a social insurer can achieve her objectives.

The LtD infrastructure comprises of introducing the new treatment with fractional coverage based random individual-level co-insurance rates. One then uses these co-insurance rates as an artificially created but an almost perfect instrumental variable to study treatment effect heterogeneity based on a spectrum of econometric tools available to researchers. Both clinical guidelines and coverage decisions can then be sequentially revised to reflect this evidence. I show that under very non-stringent conditions, the LtD infrastructure will be welfare enhancing compared to the current data production infrastructure, such as CER.

One aspect of the LtD infrastructure that would appear to be politically challenging is the notion of fractional coverage, albeit it is for a short time during the introduction of the new treatment. Although a full legal and ethical consideration of such random allocation is beyond the scope of this paper, an important note to point out is that unlike earlier discussion in this line of

reasoning that revolved around quasi-random treatment prescription (Manski 2008), the LtD infrastructure does not withhold treatment from anyone but rather changes the cost of accessing it in a random fashion. The potential for patient welfare and the richness of scientific and policy question that this infrastructure can answer should play a part in deciding its ultimate feasibility.

References

- ARROW, K.J. (1963): Uncertainty and the Welfare Economics of Medical Care. *The American Economic Review*, 53(5), 941-973.
- BASU, A. (2011): Economics of individualization in comparative effectiveness research and a basis for a patient-centered health care. *Journal of Health Economics*, 30(3), 549-559.
- BASU, A. (2014): Person-Centered Treatment (PeT) effects using instrumental variables. *National Bureau of Economic Research Working Paper No w18056. Journal of Applied Econometrics*, (In press).
- BASU, A., A.B. JENA, and T.J. PHILIPSON (2011): The impact of comparative effectiveness research on health and health care spending. *Journal of Health Economics*, 30(4), 695-706.
- DEHEJIA, R.H. (2005): Program evaluation as a decision problem. *Journal of Econometrics*, 125, 141-173.
- GARBER, A., and C. PHELPS (1997): Economic foundations of cost-effectiveness analysis. *Journal of Health Economics*, 16, 1-31.
- HECKMAN, J.J. (1996): Randomization as an instrumental variable. *The Review of Economics and Statistics*, 78(2), 336-341.
- HECKMAN, J.J. (1997): Instrumental variables: A study of implicit behavioral assumptions used in making program evaluations. *Journal of Human Resources*, 32(3), 441-462.
- HECKMAN, J.J. (2001): Accounting for heterogeneity, diversity and general equilibrium in evaluating social programmes. *The Economic Journal*, 111, F654-F699.
- HECKMAN, J.J., S. URZUA, and E. VYTLACIL (2006): Understanding instrumental variables in models with essential heterogeneity. *Review of Economics and Statistics*, 88(3), 389-432.
- HECKMAN, J.J., and E. VYTLACIL (1999): Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences*, 96(8), 4730-34.
- HECKMAN, J.J., and E. VYTLACIL (2001): *Local instrumental variables*. In C. Hsiao, K. Morimue, and J.L. Powell (Eds.) *Nonlinear Statistical Modeling: Proceedings of the Thirteenth International Symposium in Economic Theory and Econometrics: Essays in the Honor of Takeshi Amemiya*, Cambridge University Press: New York, 1-46.

- HECKMAN, J.J., and E. VYTLACIL (2005): Structural equations, treatment effects and econometric policy evaluation. *Econometrica*, 73(3), 669-738.
- MANNING, W.G., and M. S. MARQUIS (1996): Health insurance: the trade-off between risk pooling and moral hazard. *Journal of Health Economics*, 15 (5), 609–639.
- MANSKI, C. (2000): Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics* 95, 415–442.
- MANSKI, C. (2004): Statistical treatment rules for heterogeneous populations. *Econometrica* 72, 1221–1246.
- MANSKI, C. (2009): The 2009 Lawrence R. Klein Lecture: Diversified treatment under ambiguity. *International Economic Review*, 50(4), 1013-1041.
- MELTZER, D. (1997): Accounting for future costs in medical cost-effectiveness analysis. *Journal of Health Economics*, 16, 33-64.
- MOVSAS, B., J. MOUGHAN, J. OWEN, L. R. COIA, M. J. ZELEFSKY, G. HANKS, and J. F. WILSON (2007): Who enrolls onto clinical oncology trials? A radiation Patterns Of Care Study analysis. *International Journal of Radiation Oncology*Biophysics*, 68(4), 1145-50
- PAULY, M.V. (1968): The economics of moral hazard: Comment. *The American Economic Review*, 58(3), 531-537.
- PAULY, M.V. (2008): *Adverse selection and moral hazard: Implications for health insurance markets*. In Sloan F, Kasper H (eds.), *Incentives and Choice in Health and Health Care*, MIT Press, Cambridge MA.
- PAULY, M.V., and F.E. BLAVIN. (2008) Moral hazard in insurance, value-based cost sharing, and the benefits of blissful ignoring. *Journal of Health Economics*, 27, 1407-1417.
- PHILIPSON, T. J. (1997): The evaluation of new health care technology: The labor economics of statistics. *Journal of Econometrics*, 76, 375-395.
- ROY, A.D. (1951): Some thoughts on the distribution of earnings. *Oxford Economic Papers* 3(2): 135-146.
- WEINSTEIN, M., and R. ZECKHAUSER (1973): Critical ratios and efficient allocation. *Journal of Public Economics*, 2, 147-58.