Published in final edited form as:

J Health Econ. 2015 July ; 42: 165–173. doi:10.1016/j.jhealeco.2015.04.001.

Welfare implications of learning through solicitation versus diversification in health care

Anirban Basu*

*Department of Pharmacy, Health Services and Economics, University of Washington, Seattle, and NBER, Cambridge MA. 1959 NE Pacific St, Box 357660, Seattle, WA 98195-7660 (basua@uw.edu).

Abstract

Using Roy's model of sorting behavior, I study welfare implications of learning about medical care quality through the current health care data production infrastructure that relies on solicitation of research subjects. Due to severe adverse-selection issues, I show that such learning could be biased and welfare decreasing. Direct diversification of treatment receipt may solve these issues but is infeasible. Unifying Manski's work on diversified treatment choice under ambiguity and Heckman's work on estimating heterogeneous treatment effects, I propose a new infrastructure based on temporary diversification of access that resolves the prior issues and can identify nuanced effect heterogeneity.

Keywords

Learning; diversification; comparative effectiveness research; economic evaluation; instrumental variables; heterogeneity

One of the fundamental challenges in health care markets is lack of information about the quality of medical care and technology (Arrow, 1963). 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) is 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 to inform other decision makers, who device 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, substantial 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 medical technologies

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

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

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

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

In this paper, using a simple Roy's model (1951) of sorting behavior, I prove that, when incremental treatment effects are heterogeneous across patients who have access to these treatments under insurance, a data production infrastructure for comparative medical quality that relies on soliciting voluntary participation of subjects fails to identify any interpretable treatment effect parameter. Therefore, evidence generated through this process fails to inform, objectively, either the individual patient on *optimal* medical care use or a social insurer on *optimal* medical care insurance coverage.⁴

Unfortunately, such a data production infrastructure is and has been the norm for CER randomized clinical trial (RCT) studies. There are many examples of such failures in the literature. For example, Ioannidis and Lau (1997) show that in human immunodeficiency virus-related trials and trials of magnesium in acute myocardial infarction, when the benefit or toxicity from a treatment varies with the baseline risk of each patient, the treatment effect may be markedly different in populations with a different representation of high- and low-risk patients. I show that such differential representation of the population in trials may be driven more fundamentally by patient and physician behaviors and therefore the problems of interpretation of trial results are systemic.

The implications of this finding are substantial. Incomplete comparative quality information generated by CER RCTs research has the potential to misguide treatment choices since *ex-ante* perception of benefits do not coincide with the *ex-post* accrual of the same, resulting in welfare losses (Basu 2011). These inefficiencies in the choice of medical products can also accentuate the inefficiencies due to moral hazard stemming from health insurance (Arrow, 1963; Pauly, 1968), translating to higher premiums and less protection against risk, in both competitive and non-competitive insurance markets (Basu 2011). In this paper, I focus on understanding why current data production infrastructure leads to incomplete information.

I begin in the next section by laying out our ideal target parameters in CER evaluations; those that we would like to obtain estimates for in order to guide treatment decisions at the individual level and policy decisions on coverage at the social level. In Sections 2 and 3, I highlight the current data production infrastructure and prove why it would produce incomplete information. I study the implications for such incompleteness on decision-making and welfare. In Section 4, I 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 econometrics literature, one based on Manski's work on treatment choice under ambiguity (Manski 2000,

³Throughout our paper, I 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) and Malani (2008) who make distinct arguments about selection in trials of experimental therapies in the presence of health insurance

⁴Note that my assertions about optimality are very general and does not depend on specific welfare functions. What I 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.

2004, 2009) that utilizes the concept of diversification of treatment 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). I show how this framework can help overcome inefficiencies in health care markets that stem from incomplete information.

I. Defining the True Population Average Effect of a Treatment

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 effects represent the benefits (net of harms) of the new treatment over the control and are 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:

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

where $I()$ is an indicator function and α_k is interpreted as the true 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 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 I 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.

II. Data Production and Incompleteness in Quality Information

A first-best scenario can be achieved under complete information, where both the insurer and the individuals are aware of the risk classes and the production function and are able to perfectly predict b_i . If individuals had full insurance they would choose treatment if $b_i > 0$. Since the insurer can fully anticipate this individual behavior, she can provide full coverage

⁵Assume for now that the marginal cost is constant.

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

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.

Under the second-best scenario, 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). To counter this, 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.

I 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. I also assume that the social insurer's goal is to maximize consumer surplus as is realized *ex post* based on individual level choices. Throughout this paper, I 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 benefit, H^0 , is given as

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

When individuals have complete information, they choose to receive the new treatment only if $\alpha_k \geq r \cdot p$. 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 (Manning and Marquis, 1996). Consequently, the moral hazard (welfare loss) under optimal coinsurance rate (r^*) in a second-best scenario is given as

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

Eq. 4

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. $r^* \cdot p < \alpha_k < p$

Reality, however, deviates from both the first and second best scenarios, because both individuals and the social decision maker face incomplete comparative information. To understand this incompleteness, one must study the data production mechanisms in place. I consider and compare the circumstances before and after a CER study. I 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 (2009 American Recovery and Reinvestment Act; Patient Protection and Affordable Care Act of 2009, H.R. 3590, 111th Congress §630 (2010)), are called for. I 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 a CER RCT is conducted, it is safe to assume that α_k is not known with certainty at both 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 represents the heterogeneity and is the standard deviation of the effect in the population. Let individual beliefs, α'_i , be given as a single draw 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, which is based on aggregation of individual's perceived marginal benefits. 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,

$$\bar{\alpha} - p \quad \text{Eq. 5}$$

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. $r^* = 0$.

⁷I 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.

⁸This 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.

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, two groups of individuals making inefficient choices drive the welfare loss. The first group consists of people who *fail to receive treatment* because their $\alpha'_i < 0$ but they belong to risk groups 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 *receive treatment* as they are led to believe that they would get a positive benefit ($\alpha'_i \geq 0$) but realize a benefit less than the price (i.e. $\alpha_k > p$). Therefore total welfare loss is given by:

$$\begin{aligned} L_{PRE,CER} &= \left(\sum_i \sum_k I(\alpha'_i < 0) \cdot I(\alpha_k - p > 0) \cdot (\alpha_k - p) \cdot I(\Omega_i = k) \right) \\ &\quad + \left(\sum_i \sum_k I(\alpha'_i \geq 0) \cdot I(\alpha_k - p < 0) \cdot (p - \alpha_k) \cdot I(\Omega_i = k) \right) \\ &= \sum_{\alpha_k = p}^{\infty} \left\{ 1 - \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (\alpha_k - p) \\ &\quad + \sum_{\alpha_k = -\infty}^p \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (p - \alpha_k), \end{aligned} \tag{Eq. 6}$$

where $\Phi(\cdot)$ is a cumulative normal distribution. Comparing (6) to (4),

$$\begin{aligned} L_{PRE,CER}(0) &- L_{2nd}(0) = \sum_{\alpha_k = p}^{\infty} \left\{ 1 - \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (\alpha_k - p) \\ &\quad + \sum_{\alpha_k = 0}^p \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) - 1 \right\} \cdot N_k \cdot (p - \alpha_k) \\ &\quad + \sum_{\alpha_k = -\infty}^0 \left\{ \Phi\left(\frac{\bar{\alpha}}{S}\right) \right\} \cdot N_k \cdot (p - \alpha_k) \end{aligned} \tag{Eq. 7}$$

The first and the third terms in (7) are the incremental losses due to incomplete information pre-CER. The first term is the same as in (6) and comprises of individuals who fail to take treatment but would have benefited 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 positive benefits, which was not the case under the second-best scenario.

The second term in (7) 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 (7) as the first two terms drop out and $L_{PRE,CER}(0) > L_{2nd}(0)$. Thus, naturally, $L_{PRE,CER}(0) > L_{2nd}(r^*)$ for $\forall r^* \in [0,1]$, since $L_{2nd}(0) > L_{2nd}(r^*)$ 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, increase in sample size cannot overcome the issues about selection into trial that we discuss later. Nevertheless, for now,

let's assume that those concerns are not relevant. Then, as $n \rightarrow \infty$, $\bar{\alpha} \xrightarrow{p} \Delta$, $s^2 \xrightarrow{p} \sigma_\Delta^2$. Consequently, following equation (6) and substituting the post-CER estimate of average effect in it, the welfare loss post CER with an 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. 8

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, even with infinite sample, is indeterminate.

An ideal role for CER would be when, even though the social insurer continue to believe in the average effect, it can enable individual beliefs, 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, welfare losses under post ideal-CER scenario will be a modification of (6) to

$$L_{POST^*} = \sum_{\alpha_k=p}^{\infty} \left\{ 1 - \Phi \left(\frac{\alpha_k}{\sigma_k} \right) \right\} \cdot N_k \cdot (\alpha_k - p) + \sum_{\alpha_k=-\infty}^p \left\{ \Phi \left(\frac{\alpha_k}{\sigma_k} \right) \right\} \cdot N_k \cdot (p - \alpha_k) \quad \text{Eq. 9}$$

⁹In 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.

Note that both the terms in L_{POST^*} in (9) are smaller than the corresponding terms in $L_{POST,CER}$ in (6). This is because, under Assumptions 1 and 2 and also assuming $\sigma_k = s \forall k$,

$$\Phi\left(\frac{\alpha_K}{\sigma_k}\right) > \Phi\left(\frac{\bar{\alpha}}{s}\right) \text{ for } \alpha_K > p \text{ as by construction } \alpha_K > \bar{\alpha} \text{ for all } \alpha_K > p, \text{ since } \bar{\alpha} - p \geq 0.$$

Similarly, $\Phi\left(\frac{\alpha_K}{\sigma_k}\right) < \Phi\left(\frac{\bar{\alpha}}{s}\right)$ for $\alpha_K < p$ as by construction $\alpha_K < \bar{\alpha}$ for all $\alpha_K < p$. Therefore, $L_{POST^*}(r^*) < L_{PRE}(r^*) \forall r^* \in [0,1]$.

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 $\alpha_k > 0$

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 such as the Affordable Care Act of 2009 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. I will refer to such a mechanism as a randomized clinical trial (RCT) henceforth. I 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 into randomized trial implies that patients who anticipate that their expected incremental benefits from the new treatment, compared to the control, is *more* than the cost of acquiring the new treatment outside of RCT, would not enroll in the RCT and would avail the new treatment outside RCT. Similarly, patients anticipating a very small incremental benefit from the new treatment, even compared to the subsidized cost of treatment within the RCT, would not enroll. That leaves patients who enroll in a CER RCT to be those who anticipate that their incremental benefits must be *less*

than the cost of acquiring the new treatment outside of RCT but *more* that the subsidized cost of new treatment within an RCT.^{10, 11}

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 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 common 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, I show

Theorem 2—A CER randomized trial produces an unbiased estimate of the population average treatment effect (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 that are likely to be biased estimate of true average treatment effect.

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

$$S_i = I(U_i^* \geq 0), \quad \text{Eq. 10}$$

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

$$U_i^* = (\pi_R \cdot \alpha'_i - C_{RCT}) - (\alpha'_i - C_{OUT}) = (C_{OUT} - C_{RCT}) - (1 - \pi_R) \cdot \alpha'_i, \quad \text{if } \alpha'_i > 0 \quad \text{Eq. 11}$$

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

¹⁰Often enrollees are paid a monetary sum to compensate for their time spent participating in the RCTs.

¹¹It is not necessary for only the patient to be choosing the treatment. Even in acute settings, the physician choice may also be subject to the same selection principles where they have some anticipation of what works for individual patients. Therefore, physicians may often select to enroll patients to their trials.

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

$$\pi = Pr(S_i) = E(U_i^* \geq 0) = Pr(0 < \alpha'_i < (C_{OUT} - C_{RCT}) / (1 - \pi_R)) \quad \text{Eq. 12}$$

Therefore, only patients who anticipate positive benefits but whose magnitudes are less than the expected incremental cost of accessing treatment outside RCT would enroll.

Interestingly, when $C_{OUR} \approx C_{RCT}$, enrollment in CER can be quite difficult. In contrast, if $C_{OUR} \gg C_{RCT}$, then the probability of enrollment will approach one for everyone.

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 goal of RCT is to estimate a structurally meaningful population level parameter such as the average treatment effect (ATE) of treatment compared to the control. Instead, the target population ends up being defined by the RCT enrollees. Consequently, its target population defines the target parameter that an RCT tries to estimate. This parameter is an average effect that 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:

$$\Delta_{RCT} = \sum_k w_k \cdot \alpha_k \quad \text{Eq. 13}$$

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

$\pi_k = Pr(0 < \alpha'_i < (C_{OUT} - C_{RCT}) / (1 - \pi_R) | \Omega_i = k)$. The degree of selection in the trial determines the target parameter of an RCT. The touted internal validity of RCTs rests on obtaining a consistent estimate for this target parameter.

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 Eqs (2) & (13)).

On the contrary, if $\alpha'_i \not\perp b_i$, RCT since the weights would vary depending on which risk classes are more likely to enroll in the RCT.

In fact, if perceived benefits are positively correlated with true benefits, i.e.

$Corr(\alpha'_i, b_i) > 0$, it implies $Corr(w_k, \alpha_k) < 0$ 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, Eqs (12) and (13) suggest that the margin of individual who enroll in RCT anticipates a moderated positive magnitude of benefits from treatment. This 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

$E(\hat{\Delta}_{RCT}) \neq \alpha_k$. Next, I 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, I will assume $Corr(\alpha'_i, b_i) > 0$ even in the absence of a formal CER.

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):

$$\bar{\alpha} = \theta \cdot \bar{\alpha} + (1 - \theta) \cdot \hat{\Delta}_{RCT} \quad \text{Eq. 14}$$

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 or to withhold coverage on treatments with positive average net health benefits. This also highlights the fact that economic evaluations based on CER RCT studies can be misleading.

¹³It 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.

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

$$\alpha_i'' = \theta_i \cdot \alpha_i' + (1 - \theta_i) \cdot \hat{\Delta}_{RCT} \quad \text{Eq. 15}$$

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. Following (6), 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. 16

where $\mu_k(\alpha_i'') = E(\alpha_i'' | \Omega_i = k)$ and $s_k^2(\alpha_i'') = Var(\alpha_i'' | \Omega_i = k)$. Since $\frac{\mu_k(\alpha_i'')}{s_k(\alpha_i'')} < \frac{\alpha_i''}{s}$ 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 $\frac{\mu_k(\alpha_i'')}{s_k(\alpha_i'')} \geq \frac{\alpha_i''}{s}$ for $\alpha_k > p$, implying that

more individuals with $\alpha_k > p$ are drawn to take up the treatment, and $\frac{\mu_k(\alpha_i'')}{s_k(\alpha_i'')} < \frac{\alpha_i''}{s}$ for $\alpha_k < p$ implying that more individuals with $\alpha_k < p$ are drawn to not use the treatment, the CER infrastructure is welfare enhancing (i.e. $L_{POST,CER} < L_{PRE,CER}$).

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

As I have shown in the previous sections, under some general assumptions, the current CER-RCT 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 argued 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 pointed out that such an allocation automatically created randomized experiments, which were 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 (i.e. alternative coverage decisions and CER information), first physician decides whether to prescribe treatment for each individual. Second, given these potential allocations, the social decision maker decides on the level of coverage for treatment that would improve population health either by sustaining the individual choices or by incentivizing to alter them if needed. That is, a social decision maker studies these potential allocations of treatments under different information sets and makes the optimal coverage policy that would generate the highest population benefits driven by the individual allocation of treatments that would follow that policy. To the extent that one can combine the ideas of diversified treatment for 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. Here I propose a “Learning through Diversification” (LtD) infrastructure (Figure 1) that can potentially mimic the ideal CER designs discussed in Section II-B. The three main features of a LtD infrastructure are:

1. **Fractional Coverage** can be achieved using a technology lottery, : For *each new product*, develop a random order based on, say, birth dates so that this new product with uncertain effectiveness profile will be paid at varying levels by insurance (in continuous fashion) in the first year. That is, such coverage creates a completely stochastic distribution of co-insurance rates in the population, $F(b_i | \cdot) = F(b_j)$. 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.

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* (IV) to study the comparative effectiveness of receiving the new product versus its competitor and the heterogeneity in these effects in the population (Heckman 1996; Heckman 1999, 2001; Heckman and Vytlačil, 1999; Heckman et al., 2006, Basu 2012) without the challenges of voluntary participation in CER RCTs.
3. **Sequential Decision Making:** Based on the outcome evaluation results, fractional allocation rules can be adapted over time for specific risk groups. Fractional allocation would continue within risk groups where ambiguity persists. Optimal stopping rules for fractional coverage can be developed using Bayesian methods.

A. Key Features of the Learning through Diversification Infrastructure

Coinurance (demand price) as an instrument—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. 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. 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 who would and would not benefit from the most expensive technologies.

The target parameters in the LtD infrastructure—Meaningful and interpretable structural parameters for evaluation can be recovered using data arising out of an LtD infrastructure (Heckman and Vytlačil 1999, 2001; Heckman et al. 2006, Basu 2014).¹⁴ For example, local instrumental variable approaches can be used to identify Marginal Treatment Effect (MTE) parameters ((Heckman and Vytlačil 1999):

$$\frac{\partial E_g(Y|D, \Omega)}{\partial p} = E((Y_1 - Y_0) | \Omega, V=v) = MTE(\Omega, v), \quad \text{Eq. 17}$$

where $Y = D^* Y_1 + (1 - D)^* Y_0$ is the observed outcome, unobserved confounders are $V \sim \text{Uniform}[0,1]$ by construction and probability of treatment choice given the lottery can be represented by $p(\cdot, \Omega)$. Basu (2014) extends the LIV methods to identify Person-centered treatment (PeT) effects, which, for persons who choose treatment, follow

¹⁴Note 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.

$$E_{V|\Omega, P(\mathbb{D}), D} (Y_1 - Y_0 | \Omega, P(\mathbb{D}), D=1) = E(Y_1 - Y_0 | \Omega, V < P(\mathbb{D})) = P(\mathbb{D})^{-1} \int_0^{P(\mathbb{D})} MTE(\Omega, v) dv$$

Eq. 18

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 effect parameters, α_k (Eq. 1) or (Eq. 2) are readily obtained by average PeT effects over respective subgroups (Basu 2014).

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:

$$L_{CER} = L_{PRE, CER} + L_{POST, CER} \quad \text{Eq. 19}$$

where $L_{PRE, CER}$ and $L_{POST, CER}$ are given in Eq(7) and Eq(16) 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,

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

20

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. To obtain 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 be 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.

IV. CONCLUSIONS

Regulatory bodies often approves a new medical treatment 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 becomes available 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. Other countries, like UK and Canada, 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 of 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 a 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, generating 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 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. A full legal and ethical consideration of such random allocation is beyond the scope of this paper. However, it is important to note that unlike earlier discussions in this line of reasoning that revolved around quasi-random treatment prescription (Manski 2009), 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.

Acknowledgments

I am grateful for comments from Karl Claxton, David Meltzer, Justin Robertson and two anonymous reviewers and support from NIH research grants RC4CA155809 and R01CA155329. Opinions expressed are mine and do not reflect those of the University of Washington or the NBER.

References

- ARROW KJ. Uncertainty and the Welfare Economics of Medical Care. *The American Economic Review*. 1963; 53(5):941–973.
- BASU A. Economics of individualization in comparative effectiveness research and a basis for a patient-centered health care. *Journal of Health Economics*. 2011; 30(3):549–559. [PubMed: 21601299]
- BASU A. Person-Centered Treatment (PeT) effects using instrumental variables. National Bureau of Economic Research Working Paper No w18056. *Journal of Applied Econometrics*. 2014 In press.
- BASU A, JENA AB, PHILIPSON TJ. The impact of comparative effectiveness research on health and health care spending. *Journal of Health Economics*. 2011; 30(4):695–706. [PubMed: 21696840]
- DEHEJIA RH. Program evaluation as a decision problem. *Journal of Econometrics*. 2005; 125:141–173.
- GARBER A, PHELPS C. Economic foundations of cost-effectiveness analysis. *Journal of Health Economics*. 1997; 16:1–31. [PubMed: 10167341]
- HECKMAN JJ. Randomization as an instrumental variable. *The Review of Economics and Statistics*. 1996; 78(2):336–341.
- HECKMAN JJ. Instrumental variables: A study of implicit behavioral assumptions used in making program evaluations. *Journal of Human Resources*. 1997; 32(3):441–462.
- HECKMAN JJ. Accounting for heterogeneity, diversity and general equilibrium in evaluating social programmes. *The Economic Journal*. 2001; 111:F654–F699.

- HECKMAN JJ, URZUA S, VYTLACIL E. Understanding instrumental variables in models with essential heterogeneity. *Review of Economics and Statistics*. 2006; 88(3):389–432.
- HECKMAN JJ, VYTLACIL E. Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences*. 1999; 96(8): 4730–34.
- HECKMAN, JJ., VYTLACIL, E. Local instrumental variables.. In: Hsiao, C.Morimue, K., Powell, JL., editors. *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: 2001. p. 1-46.
- HECKMAN JJ, VYTLACIL E. Structural equations, treatment effects and econometric policy evaluation. *Econometrica*. 2005; 73(3):669–738.
- IONNAIDIS JP, LAU J. The impact of high-risk patients on the results of clinical trials. *Journal of Clinical Epidemiology*. 1997; 50(10):1089–98. [PubMed: 9368516]
- MALANI A. Patient enrollment in medical trials: Selection bias in a randomized experiment. *Journal of Econometrics*. 2008; 144:341–351.
- MANNING WG, MARQUIS MS. Health insurance: the trade-off between risk pooling and moral hazard. *Journal of Health Economics*. 1996; 15(5):609–639. [PubMed: 10164045]
- MANSKI C. Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics*. 2000; 95:415–442.
- MANSKI C. Statistical treatment rules for heterogeneous populations. *Econometrica*. 2004; 72:1221–1246.
- MANSKI C. The 2009 Lawrence R. Klein Lecture: Diversified treatment under ambiguity. *International Economic Review*. 2009; 50(4):1013–1041.
- MELTZER D. Accounting for future costs in medical cost-effectiveness analysis. *Journal of Health Economics*. 1997; 16:33–64. [PubMed: 10167344]
- MOVSAS B, MOUGHAN J, OWEN J, COIA LR, ZELEFSKY MJ, HANKS G, WILSON JF. Who enrolls onto clinical oncology trials? A radiation Patterns Of Care Study analysis. *International Journal of Radiation Oncology*Biophysics*. 2007; 68(4):1145–50.
- PAULY MV. The economics of moral hazard: Comment. *The American Economic Review*. 1968; 58(3):531–537.
- PAULY, MV. Adverse selection and moral hazard: Implications for health insurance markets.. In: Sloan, F., Kasper, H., editors. *Incentives and Choice in Health and Health Care*. MIT Press; Cambridge MA: 2008.
- PAULY MV, BLAVIN FE. Moral hazard in insurance, value-based cost sharing, and the benefits of blissful ignoring. *Journal of Health Economics*. 2008; 27:1407–1417. [PubMed: 18760850]
- PHILIPSON TJ. The evaluation of new health care technology: The labor economics of statistics. *Journal of Econometrics*. 1997; 76:375–395.
- ROY AD. Some thoughts on the distribution of earnings. *Oxford Economic Papers*. 1951; 3(2):135–146.
- WEINSTEIN M, ZECKHAUSER R. Critical ratios and efficient allocation. *Journal of Public Economics*. 1973; 2:147–58.

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