

NBER WORKING PAPER SERIES

SELECTIVE TRIALS:
A PRINCIPAL-AGENT APPROACH TO RANDOMIZED CONTROLLED EXPERIMENTS

Sylvain Chassang
Gerard Padro i Miquel
Erik Snowberg

Working Paper 16343
<http://www.nber.org/papers/w16343>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2010

We are particularly indebted to Abhijit Banerjee, Roland Benabou, and Jeff Ely for advice and encouragement. The paper benefited greatly from conversations with Attila Ambrus, Nava Ashraf, Oriana Bandiera, Angus Deaton, Esther Duflo, Pascaline Dupas, Greg Fischer, Kripa Freitas, Drew Fudenberg, Justin Grimmer, Rema Hanna, Jim Heckman, Johannes Hörner, Dean Karlan, Michael Kremer, Guido Imbens, John Ledyard, Maggie McConnell, Stephen Morris, Muriel Niederle, Marcin Peski, Nancy Qian, Antonio Rangel, Imran Rasul, Dan Scharfstein, Sam Schulhofer-Wohl, Jesse Shapiro, Monica Singhal, Andy Skrzypacz, Francesco Sobbrío, Lars Stole, Chris Woodruff and Eric Zitzewitz, as well as seminar participants at Boston University, Brown, Caltech, Chicago Booth, Harvard/MIT, HEC Lausanne, Johns Hopkins, LSE, MPSA, NYU Stern, the Radcliffe Institute, Stanford, Stockholm School of Economics, SWET, UT Austin, Washington University in St. Louis, and Yale. Part of this work was done while Chassang visited the Department of Economics at Harvard, and he gratefully acknowledges their hospitality. Paul Scott provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2010 by Sylvain Chassang, Gerard Padro i Miquel, and Erik Snowberg. 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.

Selective Trials: A Principal-Agent Approach to Randomized Controlled Experiments
Sylvain Chassang, Gerard Padro i Miquel, and Erik Snowberg
NBER Working Paper No. 16343
September 2010
JEL No. C9,C90,C93,D82,O12,O22

ABSTRACT

We study the design of randomized controlled experiments in environments where outcomes are significantly affected by unobserved effort decisions taken by the subjects (agents). While standard randomized controlled trials (RCTs) are internally consistent, the unobservability of effort provision compromises external validity. We approach trial design as a principal-agent problem and show that natural extensions of RCTs -which we call selective trials- can help improve the external validity of experiments. In particular, selective trials can disentangle the effects of treatment, effort, and the interaction of treatment and effort. Moreover, they can help experimenters identify when measured treatment effects are affected by erroneous beliefs and inappropriate effort provision.

Sylvain Chassang
Department of Economics
Bendheim Hall 316
Princeton University
Princeton, NJ 08544-1021
chassang@princeton.edu

Erik Snowberg
Division of the Humanities
and Social Sciences 228-77
California Institute of Technology
Pasadena, CA 91125
snowberg@caltech.edu

Gerard Padro i Miquel
STICERD
London School of Economics
Houghton Street
London, WC2A 2AE
United Kingdom
and NBER
g.padro@lse.ac.uk

1 Introduction

This paper studies the design of experimental trials when outcomes depend significantly on unobserved effort decisions taken by subjects (agents).¹ Even in an ideal setting where the experimenter (principal) can randomly and independently assign an arbitrarily large number of agents to the treatment and control groups, unobserved effort limits the informativeness of randomized controlled trials (RCTs). For example, if a technology’s measured returns are low, it is difficult to distinguish whether this occurred because true returns *are low* or because most agents *believe they are low* and therefore put no effort into using the technology. Moreover, to the extent that effort responds to beliefs, and beliefs respond to information, this makes it difficult to predict the returns to the technology on the same population as it becomes better informed. In other words, unobserved effort is a source of heterogeneity in treatment effects, and is a significant challenge to the external validity of experimental trials.²

We propose simple extensions to RCTs—which we call selective trials—that improve the external validity of trial results without sacrificing robustness or internal validity. These experimental designs can be used to determine the extent to which inappropriate effort or erroneous beliefs affect treatment effects. We provide a systematic analysis of trial design using a principal-agent framework with both adverse selection—an agent’s type is unobserved—and moral hazard—an agent’s effort is unobserved. However, unlike the standard principal-agent framework, our principal’s goal is to maximize information about a technology’s returns (in the sense of Blackwell) rather than profits. The principal seeks to achieve this objective through single-agent mechanisms that assign agents to treatments of varying sophistication based on the message they send.

¹Throughout the paper we call experimental subjects agents, and call the experimenter the principal. Following usual conventions, we refer to the principal as she and an agent as he.

²Unobserved effort is an issue whether a trial is open—agents know their treatment status—or blinded—agents’ treatment status is obscured by giving the control group a placebo. See Duflo et al. (2008b) for a more detailed description of RCTs and the external validity issues frequently associated with them.

These mechanisms improve on RCTs for two reasons. First, they let agents express preferences over their treatment by probabilistically selecting themselves in and out of the treatment group at a cost (hence the name selective trials).³ This makes implicit, unobserved selection an explicit part of the experimental design. Second, these mechanisms allow for treatments of varying richness: in open trials, treatment corresponds to access to the new technology; in blind trials, treatment corresponds to an undisclosed allotment of the technology, as well as information over the likelihood of having been allotted the technology; and in incentivized trials, treatment corresponds to access to the technology as well as an incentive (or insurance) contract based on outcomes.

Our results fall into two broad categories. Given a type of treatment (open, blind or incentivized), our first set of results establishes conditions under which a large sample mechanism is maximally informative and examines the small sample properties of such mechanisms. We show that a mechanism is maximally informative if and only if it identifies an agent's preferences over all possible treatment assignments and, given preferences, still assigns each agent to the treatment or control group with positive probability. Thus, our designs encapsulate the data generated by a standard randomized controlled trial. These designs can be implemented in a number of intuitive ways, such as a menu of lotteries or utilizing the design of Becker et al. (1964), referred to as the BDM mechanism.

In small samples, selective trials have some costs because any mechanism that identifies agents' preferences in a strictly incentive compatible way must assign agents with a higher value for the technology to the treatment group with higher probability. This oversampling of high value agents is an additional constraint which can reduce power. However, these sampling costs can be reduced by weakening incentives for truthfully reporting preferences, so the experimenter can strike a balance between sampling costs and the precision of the preference data that is obtained. As we detail later, these results contribute to recent

³For simplicity, we focus on monetary costs, but the mechanisms can be based on non-monetary costs. For example, agents could choose between lines of different lengths to place themselves into the treatment group with different probabilities.

discussions over the usefulness of charging subjects for access to treatment in RCTs (see for instance Cohen and Dupas (2010), Dupas (2009b), or Ashraf et al. (forthcoming)).

Our second class of results characterizes what can be inferred from selective trials, and highlights how they contribute to the ongoing discussion on the external validity of field experiments (Deaton, 2010; Imbens, 2010).⁴ By eliciting agents' value for the technology, open selective trials recover the distribution of returns as a function of willingness to pay. As a result, open trials provide a simple and robust way to recover the marginal treatment effects (MTEs) introduced by Heckman and Vytlacil (2005). While MTEs can be used to extrapolate the treatment effect of policies affecting the accessibility of goods, such as subsidies, they do not typically allow projections about interventions that alter beliefs and effort provision, such as informational campaigns.

Selective trials go beyond MTEs and identify deep parameters by letting agents express preferences over richer treatments. Specifically, we consider blind trials where treatment status is hidden from agents by giving the control group a placebo. This allows us to vary the information an agent has over his treatment status. As a result we can identify the pure effect of treatment, as well as the agents' real and perceived returns to effort.⁵ As blind trials

⁴In addition, selective trials may alleviate subversions of experimental protocol discussed in Deaton (2010). That is, explicitly allowing the agents to select themselves in and out of treatment may reduce the number of agents in the control group who obtain the treatment by other means, as well as the number of agents in the treatment group that refuse to be treated. Furthermore, the principal may use the information revealed by agents' preferences to increase monitoring of agents who expressed a high value for treatment but were assigned to the control group.

Note that the percentage of agents rejecting, or opting-in to, treatment is often significant. For example, 45% of the people Dupas and Robinson (2009) opened a savings account for never made a deposit, 72% of the people offered a commitment saving product by Ashraf et al. (2006) rejected it, and in a study of educational vouchers in Columbia, Angrist et al. (2002) find that 25% of those randomly denied a voucher were awarded other scholarships, and 10% of those who were offered vouchers declined them.

⁵Although uncommon in economics, blind trials are quite common in medicine. For a brief review of RCTs in medicine see Stolberg et al. (2004). Jadad and Enkin (2007) provides a more comprehensive review. Selective trials nest preference trials, which have generally been used in medicine to assess the ethics of using randomized controlled trials. A common implementation of preference trials compares outcomes from a standard randomized trial to results from a trial in which agents can perfectly select their treatment status. This provides information about whether or not, in a given environment, letting subjects choose their preferred treatment confounds the evaluation of treatment effects. Our work shows that eliciting preferences is not incompatible with randomization, and that preferences carry information that facilitates inference from treatment effects. For more on preference trials, see Zelen (1979); Flood et al. (1996); Silverman and

are rarely used in economics—often for want of a convincing, ethical placebo—we extend the analysis to incentivized trials in which agents are informed of their treatment status, but receive different transfers conditional on observable outcomes. Under mild assumptions, this produces information similar to that produced by selective blind trials. While the experimental designs we propose may pose some implementation challenges, many elements of selective trials have already been used successfully in field studies (see, for example, Ashraf et al., forthcoming; Karlan and Zinman, 2010; Cohen and Dupas, 2010; Berry et al., 2010), and we believe that the designs we suggest can be gainfully applied in practice.

The literature on treatment effects, based on a statistical framework quite different from our principal-agent approach, has largely focused on much simpler effort decisions and the ex post analysis of data. In this literature, agents are usually viewed as either taking treatment or not (with the notable exceptions of Angrist and Imbens (1995) and Jin and Rubin (2008)), and more importantly, this decision is assumed to be observable (or sufficiently correlated with exogenous observable variables) and based on correct beliefs about returns (Imbens and Angrist, 1994; Angrist et al., 1996; Heckman and Vytlačil, 2005). In contrast, we consider effort decisions which are unobservable, high dimensional, and can be the result of incorrect beliefs. Additionally, most previous approaches, even those which rely—as we do—on decision theory, focus on modeling data from an RCT after it has been run (Philipson and Desimone, 1997; Philipson and Hedges, 1998).⁶ We take an ex ante perspective and propose designs for experimental trials that can help understand how beliefs and effort affect treatment effects.

The paper is organized as follows. Section 2 uses a simple example to illustrate the main points of the paper. Section 3 defines the general framework. Section 4 investigates selective open trials. Section 5 turns to blind selective trials and shows how they can be used to identify true and perceived returns to effort. Section 6 extends the analysis to incentivized

Altman (1996); King et al. (2005); Jadad and Enkin (2007); Tilbrook (2008).

⁶There is a large literature on experimental design that considers issues that we largely take for granted, such as the efficient implementation of randomization.

trials and shows that under reasonable assumptions they can be as informative as blind selective trials, without placebos. Section 7 concludes with a discussion of the limitations of and future directions for our approach to designing randomized controlled experiments.

2 An Example

This section uses a highly stylized example to illustrate the paper’s main points. While the model in this section is simple, the notation and concepts will carry through to the more general framework unless specifically noted. To fix ideas, we use the example of an experiment evaluating the health effects of a water treatment product.⁷

2.1 A Simple Model

There are infinitely many agents indexed by $i \in \mathbb{N}$. Each agent has a treatment status $\tau_i \in \{0, 1\}$. If agent i is in the treatment group, $\tau_i = 1$, and he is given the water treatment product. Otherwise $\tau_i = 0$ and the agent is in the control group.

Agent i obtains a final outcome $y_i \in \{0, 1\}$, which can be measured by the principal. In our example $y_i = 1$ indicates that the agent has remained healthy. The probability that an agent remains healthy depends on both treatment and effort:

$$\text{Prob}(y_i = 1 | e_i, \tau_i) = q_0 + R e_i \tau_i \tag{1}$$

where $e_i \in [0, 1]$ is agent i ’s decision of whether or not to put effort into using the product, $R \in [R_L, R_H]$ is the component of the technology’s return that is common to all agents and

⁷It should be noted that while our main focus is on the use of RCTs in medical, public health and development contexts, our analysis applies to most environments involving decentralized experimentation. For instance, if a firm wants to try a new way to organize production, specific plant managers will have to decide how much effort to put towards implementing it. The firm’s CEO is in the same position as the principal in our framework, and must guess the effort exerted by his managers when evaluating returns to the new production scheme. Similarly, if a school board wants to experiment with a new program, individual teachers and administrators will have to decide how much effort to expend on implementing the program.

q_0 is the unknown baseline likelihood of staying healthy over the study period, which will be controlled for using randomization. Agents have different types t which characterize their beliefs over returns R . We denote by $R_t = \mathbb{E}_t R$ the returns expected by an agent of type t . The distribution F_{R_t} , of expectations R_t in the population, need not be known to the principal or the agents.⁸

We assume throughout that effort is private and cannot be monitored by the principal. In other words, we assume that all observable dimensions of effort are already controlled for, and focus on those dimensions that are not observable. For example, with a water treatment product, an experimenter may be able to determine whether or not the agent has treated water in his home, but it may be much more difficult to determine if the agent drinks treated water when away from home.⁹

Given effort e_i , agent i 's expected utility is given by

$$\mathbb{E}_t[y_i|e_i] - ce_i, \tag{2}$$

where $c \in (R_L, R_H)$ is the agents' cost of effort.¹⁰ In our example, this may be the cost of remembering to use the product, the social cost of refusing untreated water, or disliking the taste of treated water. In addition, we assume each agent has quasilinear preferences with respect to money. An agent's willingness to pay for treatment is $V_t = \max\{R_t - c, 0\}$, which we assume is less than some value V_{\max} for all agents.

We focus initially on open trials where agents know their treatment status before making effort decisions, and contrast two ways of running trials: a standard RCT, where agents are randomly assigned to the treatment group with probability π , and a selective open trial which

⁸We focus on heterogenous beliefs as a source of heterogenous behavior and heterogenous returns as, in this setting, convincingly identifying true returns to treatment would be particularly valuable and have a large effect on behavior. The general framework, described in Section 3, allows for general, idiosyncratic, returns.

⁹Still, as Duflo et al. (2010) shows, innovative monitoring technologies can be quite useful. To the extent that monitoring is possible, it should be done.

¹⁰In this example, allowing c to vary with type does not change any of the results.

lets agents express preferences over treatment by selecting their probability of treatment.

We implement a selective trial here using the BDM mechanism:

- The agent sends a message $m \in [0, V_{\max}]$ indicating his willingness to pay for treatment.
- A price p to obtain treatment is drawn according to a distribution with convex support, and c.d.f. F_p such that $0 < F_p(0) < F_p(V_{\max}) < 1$.
- If $m \geq p$, the agent obtains the treatment at price p , otherwise, the agent is in the control group and no transfers are made.

Note that a higher message m , increases an agent's probability of treatment, $F_p(m)$, as well as his expected payment: $\int_{p \leq m} p dF_p$.

2.2 The Limits of RCTs and the Value of Self-Selection

Inference from Randomized Controlled Trials. We begin by considering the information produced by an RCT. If agent i is in the treatment group, he chooses to expend effort ($e = 1$) if and only if $R_t \geq c$. Hence, the average treatment effect identified by an RCT is¹¹

$$\begin{aligned} \Delta^{RCT} &= \mathbb{E}[y|\tau = 1] - \mathbb{E}[y|\tau = 0] \\ &= \mathbb{E}[q_0 + R \times \mathbf{1}_{R_t \geq c} | \tau = 1] - \mathbb{E}[q_0 | \tau = 0] \\ &= R \times \text{Prob}(R_t > c) = R \times (1 - F_{R_t}(c)). \end{aligned}$$

When the distribution of agents' expectations F_{R_t} is known, then an RCT will identify R . However, in most cases F_{R_t} is not known, and average treatment effect Δ^{RCT} provides a garbled signal of the underlying returns R . If the outcomes of agents in the treatment group

¹¹In the medical literature, R is referred to as the *efficacy* of a treatment, and Δ^{RCT} , which identifies the average treatment effect, is referred to as the *effectiveness* of the treatment. While effectiveness varies with the beliefs and effort decisions of agents in the experimental population, efficacy does not. Moreover, R is similar to the complier average causal effect (CACE) which Imbens and Angrist (1994) and Angrist et al. (1996) have shown is identified if effort is observable and can only take on one of two values.

are not particularly good compared to agents in the control group, the principal does not know if this is because the water treatment product is not particularly useful, or because the agents did not put sufficient effort towards using the treatment.

Inference from Open Selective Trials. We now turn to selective trials and show they are more informative than RCTs.

The selective trial described above elicits agents’ willingness to pay and, conditional on a given willingness to pay V , generates non-empty treatment and control groups. As F_p has convex support, it is a strictly dominant strategy for an agent to submit a message $m = V_t$ equal to his willingness to pay for treatment. Thus, an agent with value V_t has probability $F_p(V_t)$ of being in the treatment group and probability $1 - F_p(V_t)$ of being in the control group. Both of these quantities are strictly positive as $0 < F_p(0) < F_p(V_{\max}) < 1$.¹²

The selective trial described above provides us with the set of local instruments needed by Heckman and Vytlačil (2005) to estimate marginal treatment effects (MTEs). That is, for any willingness to pay V , we are able to estimate,

$$\begin{aligned} \Delta^{MTE}(V) &\equiv \mathbb{E}[y|\tau = 1, V_t = V] - \mathbb{E}[y|\tau = 0, V_t = V] \\ &= \mathbb{E}[y|\tau = 1, m_t = V] - \mathbb{E}[y|\tau = 0, m_t = V] \end{aligned}$$

which can be used to perform policy simulations in which the distribution of types is constant but access to the technology is changed—for example, subsidies. Moreover, MTEs can be integrated to recover the average treatment effect identified by an RCT.

In the current environment, because willingness to pay is a good signal of future use, MTEs can be used to identify the true returns R . Specifically, all agents with value $V_t > 0$ also believe $R_t - c > 0$ and hence, put effort $e = 1$ towards using the technology.¹³ Hence,

¹²Note also that agents with higher value are treated with higher probability. This matters for the precision of estimates in small samples, a point we return to in Section 4.

¹³Note that in this very simple environment the same result is obtained by setting a price p and selecting a probability of randomization π such that $F_{R_t}(\frac{p}{\pi} - c)$, and then examining the treatment effect for those that

it follows that

$$\begin{aligned}\Delta^{MTE}(V > 0) &= \mathbb{E}[q_0 + R \times e_t | \tau = 1, V_t > 0] - \mathbb{E}[q_0 | \tau = 0, V_t > 0] \\ &= R.\end{aligned}$$

A selective trial identifies the average treatment effect, MTEs, and true returns R . Hence, it is more informative than an RCT, which only identifies the average treatment effect.

The true return R and the distribution of valuations V_t have several policy uses. First, knowing R allows us to simulate the treatment effect for a population where everyone puts in the appropriate amount of effort. Second, these variables allow us to estimate the returns to increasing usage within a given population. Third, and finally, the data provided by the selective trial can be used to inform agents and disrupt learning traps more effectively than data from an RCT. For example, imagine that the true returns to the technology are high, but most agents believe they are low. In that case, an RCT will measure low returns to the treatment and will not convince agents that they should be expending more effort. In contrast, the data generated by a selective trial would identify that true returns are high, lead agents to update their beliefs, and efficiently adopt the water treatment product.¹⁴

2.3 Richer Treatments

In the previous subsection, a selective trial identified true returns because willingness to pay was a good predictor of future usage. However, as our continuing example shows, this will not always be the case. Thus, MTEs are generally not sufficient to infer whether beliefs are affecting measured treatment effects. However, more sophisticated selective trials such as blind selective trials and incentivized selective trials can be used to recover true returns.

pay the price p . The idea that a higher price will select individuals who will use a product more intensely has been in the economics literature for some time and is closely related to classic selection models. See Roy (1951) and Oster (1995).

¹⁴For more on the effect of appropriate information on behavior, see Thornton (2008), Dupas (2009a) or Nguyen (2009).

We modify the example so that the returns R to the technology include both a baseline return and returns to effort: $R = (R_b, R_e) \in \mathbb{R}^2$. In the context of a water treatment product, R_b could be the baseline returns to using the water treatment product only when it is convenient to do so, and R_e the additional returns to using it more thoroughly (for example, bringing treated water when away from home). Success rates given effort and treatment status are:

$$\begin{aligned}\text{Prob}(y = 1|\tau = 0, e) &= q_0 \\ \text{Prob}(y = 1|\tau = 1, e) &= q_0 + R_b + eR_e.\end{aligned}$$

An agent of type t has expectation $(R_{b,t}, R_{e,t})$ over returns $R = (R_b, R_e)$, and expends effort if and only if $R_{e,t} \geq c$. Therefore, an agent's willingness to pay for treatment is given by $V_t = R_{b,t} + \max\{R_{e,t} - c, 0\}$.

Inference from Open Selective Trials. We have already shown that open selective trials can identify treatment effects conditional on willingness to pay. However, in the current environment, willingness to pay is no longer a good signal of effort. Indeed, there are now two reasons why an agent might value the treatment: he believes that a thorough use of the product has high returns ($R_{e,t}$ is high)—the channel emphasized in Section 2.2—or he believes that a casual use of the water treatment product is sufficient to obtain high returns and that thorough use brings little additional return ($R_{b,t}$ is high, but $R_{e,t}$ is low). Hence, agents who are willing to pay because they think baseline returns are high need not be the agents who will actually expend effort. Formally, a selective trial still identifies MTEs,

$$\Delta^{MTE}(V) = R_b + R_e \text{Prob}[R_{e,t} \geq c | R_{b,t} + \max\{R_{e,t} - c, 0\} = V],$$

but these are generally not sufficient to recover R_b and R_e .¹⁵ As a result, MTEs are insufficient to simulate the returns of a population of agents that all expended appropriate effort, or more generally, the returns to increasing the effort of agents. Nor do MTEs provide the information needed for the agents to infer true returns.

Blind Selective Trials. In a blind trial, the agent does not know his treatment status $\tau \in \{0, 1\}$ at the time of effort, but rather knows his probability $\phi \in [0, 1]$ of having been assigned to the treatment group. Open trials are blind trials where ϕ is either 0 or 1.

Given a probability ϕ of being treated, the agent puts effort if and only if $\phi R_{e,t} - c > 0$. The agent's expected value for being treated with probability ϕ is

$$V_t(\phi) = \phi R_{b,t} + \max\{\phi R_{e,t} - c, 0\}.$$

We depart from standard blind trials in a simple but fundamental way: while standard blind trials keep ϕ fixed and do not infer anything from the specific value of ϕ used, we allow ϕ to vary and use both willingness to pay and outcomes at different values of ϕ for inference.

As with open trials, willingness to pay can be elicited using a BDM-type mechanism. However, as willingness to pay, $V_t(\phi)$, now depends on ϕ , the mechanism in Section 2.1 is implemented after the agent is asked to send a message $m(\phi)$ for each possible value of ϕ . A value of ϕ is drawn from a c.d.f. F_ϕ , which has mass points at 0 and 1, and p is independently drawn from a c.d.f. F_p , as before. If $m(\phi) \geq p$, the agent pays p and is allotted the treatment with probability ϕ ; otherwise, the agent is in the control group and no transfers are made.

A first advantage of blind trials is that, unlike open trials, an agent's actual treatment status τ and his belief over his treatment status are different. This allows for a robust identification of baseline returns R_b . If an agent is assigned a probability of treatment $\phi > 0$

¹⁵For instance, it is not possible to distinguish a situation in which returns to effort are equal to R_e and a proportion ηV of agents with value V puts effort, from a situation in which returns to effort are $2R_e$ and a proportion $\frac{\eta}{2}V$ of agents with value V puts effort.

low enough that $\phi R_H < c$, he will not expend any effort. Still, a proportion $\phi > 0$ of these agents do receive treatment while a proportion $1 - \phi > 0$ do not. Hence we can identify R_b by measuring the effect of treatment for agents known not to exert effort:

$$R_b = \mathbb{E} \left[y \mid \phi < \frac{c}{R_H}, \tau = 1 \right] - \mathbb{E} \left[y \mid \phi < \frac{c}{R_H}, \tau = 0 \right].$$

A second advantage of blind trials is that the agents' value mapping $V_t(\phi)$ allows identification of which agents expend effort when treated for sure. Let $\theta_t \equiv \frac{1}{2}V_t(\phi=1) - V_t(\phi=1/2)$. Given that $V_t(\phi=0) = 0$, θ_t is the value that an agent with belief $\phi = 1/2$ is willing to pay to learn his treatment status. If an agent with belief $\phi = 1$ expends effort then

$$\begin{aligned} \theta_t &= \frac{1}{2}[R_{b,t} + R_{e,t} - c] - \frac{1}{2}R_{b,t} - \max \left\{ \frac{1}{2}R_{e,t} - c, 0 \right\} \\ &\geq \min \left\{ \frac{R_{e,t} - c}{2}, \frac{c}{2} \right\} > 0, \end{aligned}$$

i.e. the agent has positive willingness to pay for information. Conversely, if an agent does not intend to put effort when $\phi = 1$ then there is no value for information, and $\theta_t = 0$. Hence, the sign of θ_t provides a simple test for whether or not the agent would put effort given treatment. Since F_ϕ puts positive mass at $\phi = 0$ and $\phi = 1$ and $F_p(0) > 0$, given any value function, a positive mass of agents get $\phi = 0$ and a positive mass of agents get $\phi = 1$. Thus, provided that some agents satisfy $\theta_t > 0$, we can identify R_e using either of the following expressions:

$$\begin{aligned} R_e &= \mathbb{E}[y \mid \phi = 1, \theta_t > 0, \tau = 1] - \mathbb{E}[y \mid \phi = 1, \theta_t = 0, \tau = 1] \\ &= \mathbb{E}[y \mid \phi = 1, \theta_t > 0, \tau = 1] - \mathbb{E} \left[y \mid \phi < \frac{c}{R_H}, \theta_t > 0, \tau = 1 \right]. \end{aligned}$$

Incentivized Selective Trials. We now show that incentivized trials can provide the experimenter with similar information to blind trials. This is useful as in many areas of eco-

nomic interest, blind trials are not practical due to the lack of suitable, or ethical, placebos.

In an incentivized selective trial, the agent obtains a treatment status $\tau \in \{0, 1\}$, makes a fixed transfer p (which can be positive or negative), and receives a bonus (or penalty) w in the event that $y = 1$. Note that if $p > 0$ and $w > 0$, then the agent is being assigned an incentive contract. If instead $p < 0$ and $w < 0$, the agent is assigned an insurance contract.

Given a bonus level w , the agent puts effort if and only if $(1 + w)R_{e,t} - c > 0$. In turn, the agent's willingness to pay for treatment given bonus w is

$$V_t(w) = (1 + w)R_{b,t} + \max\{(1 + w)R_{e,t} - c, 0\}.$$

As before, the mapping $w \mapsto V_t(w)$ can be elicited using a variant of the BDM mechanism. Incentivized trials allow us to evaluate baseline returns in a straightforward manner. When offered a full insurance contract $w = -1$, the agent will put effort $e = 0$ so that

$$R_b = \mathbb{E}[y|w = -1, \tau = 1] - \mathbb{E}[y|w = -1, \tau = 0].$$

In turn, notice that for any type t with $R_{e,t} > 0$, there exists a value w_t such that whenever $w > w_t$, the agent expends effort $e = 1$. Value w_t is identified from mapping $w \mapsto V_t(w)$ since

$$\left. \frac{\partial V_t}{\partial w} \right|_{w > w_t} = R_{e,t} + R_{b,t} > R_{b,t} = \left. \frac{\partial V_t}{\partial w} \right|_{w < w_t}.$$

Additionally, this last expression allows us to identify the agent's subjective beliefs over baseline returns and returns to effort $(R_{b,t}, R_{e,t})$. Pick some value \bar{w} sufficiently high that it induces some agents to put effort, and construct statistic $\bar{w} - w_t$. Returns to effort can be identified by

$$\begin{aligned} R_e &= \mathbb{E}[y|w = \bar{w}, \bar{w} - w_t > 0, \tau = 1] - \mathbb{E}[y|w = \bar{w}, \bar{w} - w_t < 0, \tau = 1] \\ &= \mathbb{E}[y|w = \bar{w}, \bar{w} - w_t > 0, \tau = 1] - \mathbb{E}[y|w = -1, \bar{w} - w_t > 0, \tau = 1]. \end{aligned}$$

Just like blind trials, incentivized trials identify true returns $R = (R_b, R_e)$.

Altogether, this section suggests that while unobserved effort is an issue for the external validity of standard randomized controlled trials, appropriate ex ante trial design—rather than ex post data treatment—may help in alleviating these concerns. However, these results are obtained using a particularly simple framework, and their robustness must be investigated. The remainder of the paper extends the analysis to a very general framework to provide systematic results about which mechanisms are the most informative, what their small sample properties are, and what can be inferred from the data they generate.

3 A General Framework

We now generalize the framework used in our example. Once again, there are infinitely many agents, indexed by $i \in \mathbb{N}$.¹⁶ Returns to the technology are described by parameter $R \in \mathcal{R} \subset \mathbb{R}^\kappa$.

Types. Each agent i has a type $t \in T$, which includes a belief over returns R , as well as factors that might affect behavior and outcomes, such as idiosyncratic costs of effort, idiosyncratic returns, and beliefs over such factors. We assume that agents are exchangeable, so that their types are i.i.d. draws from some distribution $\chi \in \Delta(T)$, which is itself a random variable. A profile of types is given by $\mathbf{t} \in T^\mathbb{N}$. For concision we omit publicly observable traits, but it is straightforward to allow for them.

Outcomes and Success Rates. Agent i obtains an outcome $y_i \in \{0, 1\}$.¹⁷ An agent’s true and perceived likelihoods of success (that is, $\text{Prob}(y = 1)$) depend on his type, the aggregate returns to the technology and the agent’s effort choice $e \in E$, where E is a compact subset

¹⁶We will discuss how our results change with finitely many agents.

¹⁷As Appendix A shows, this greatly simplifies notation but is not essential to our results.

of $\mathbb{R}^{k'}$.¹⁸ Success rates are denoted by

$$\begin{aligned} q(R, t, \tau_i, e_i) &= \text{Prob}(y=1|R, t, \tau_i, e_i) \\ q_t(\tau_i, e_i) &= \int_R q(R, t, \tau_i, e_i) dt(R) \end{aligned}$$

where $q(R, t, \tau, e)$ is the true success rate of an agent of type t (this allows for idiosyncratic returns) while $q_t(\tau, e)$ is the probability of success perceived by an agent of type t .¹⁹ We assume that q and q_t are continuous with respect to effort e .

Preferences. Given effort e_i , treatment status τ_i , monetary transfer p_i , and final outcome y_i , agent i 's utility is $u(y_i, t_i) - c(e_i, t_i) - p_i$.²⁰

Assignment Mechanisms. We distinguish three ways to assign treatment:

1. *Open selective trials* are mechanisms $G_o = (M_o, \mu_o)$ where M_o is a set of messages and $\mu_o : M_o \rightarrow \Delta(\{0, 1\} \times \mathbb{R})$ maps individual messages to a probability distribution over treatment status τ_i and transfers p_i .
2. *Blind selective trials* are mechanisms $G_b = (M_b, \mu_b)$ where M_b is a set of messages and $\mu_b : M_b \rightarrow \Delta([0, 1] \times \mathbb{R})$ maps messages to a probability distribution over uncertain treatment status ϕ_i (where $\phi_i = \text{Prob}(\tau_i = 1)$) and transfers p_i .
3. *Incentivized selective trials* are mechanisms $G_w = (M_w, \mu_w)$ where M_w is a set of messages and $\mu_w : M_w \rightarrow \Delta(\{0, 1\} \times \mathbb{R} \times \mathbb{R})$ maps messages to a probability distribution

¹⁸In most settings, this effort decision is multidimensional. For instance, in the case of fertilizer, it is not enough for agents to just expend effort spreading fertilizer. As Duflo et al. (2008a) highlight, effort is needed to choose the appropriate seeds to go with the fertilizer, learn how much and when to water the crops, and to learn how much fertilizer gives the highest returns at the lowest cost. In this case it is natural to think of effort as a vector, where the first component corresponds to picking the right seeds, the second to the right amount of fertilizer, the third to properly applying it, and so on.

¹⁹Note that although returns conditional on the state R are common knowledge, heterogeneous priors allow for arbitrary disagreements between the principal and the agents.

²⁰Note that p_i can be negative, or that all transfers can be rescaled by a fixed amount to improve participation. See Appendix A for a treatment of the case where agents have non-quasilinear preferences.

over treatment status τ_i , a fixed transfer p_i from the agent to the principal, and a bonus w_i transferred from the principal to the agent conditional on $y_i = 1$.

Note that these are single agent mechanisms. Agent i 's final assignment depends only on his message, and not on messages sent by others (see Section 7 for a discussion of multi-agent mechanisms). We denote by $\pi(m) = \text{Prob}(\tau = 1|m)$ the likelihood of being given the treatment when sending message m . We focus largely on mechanisms G such that χ -almost surely, every agent i has a dominant message $m_G(t_i)$. In all these designs agents can probabilistically select their assignment using messages, hence, the name *selective trials*.

Informativeness of Mechanisms. We evaluate mechanisms according to their informativeness, in the sense of Blackwell. We say that a mechanism G is at least as informative as a mechanism G' (denoted by $G' \preceq G$) if the data generated by G' can be simulated using only data generated by G .

Specifically, denote by a_i the assignment given to agent i by whichever mechanism is chosen. The principal observes data $\mathbf{d}_G = (m_i, a_i, y_i)_{i \in \mathbb{N}}$. Denote \mathcal{D}_G the set of possible data sequences generated by mechanism G . Mechanism G is at least as informative as mechanism G' , denoted by $G' \preceq G$, if and only if there exists a fixed data manipulation procedure $h : \mathcal{D}_G \rightarrow \Delta(\mathcal{D}_{G'})$ such that for all $\mathbf{t} \in T^{\mathbb{N}}, R \in \mathcal{R}$, $h(\mathbf{d}_G(\mathbf{t}, R)) \sim \mathbf{d}_{G'}(\mathbf{t}, R)$.

This notion of informativeness is easier to work with in environments with infinite samples, as this focuses on issues of identification rather than issues of statistical power. However, this definition also applies in the case of finitely many agents.

4 Open Selective Trials

In open selective trials an agent is assigned a treatment status τ and a transfer p based on message m . Given this assignment (τ, p) , the indirect utility of an agent with type t is

$V_t(\tau) - p$ where,

$$V_t(\tau) = \max_{e \in E} q_t(\tau, e)u(y=1, t) + [1 - q_t(\tau, e)]u(y=0, t) - c(e, t).$$

Since flow utility is identified up to a constant, we can normalize the value of being in the control group $V_t(\tau = 0)$ to zero for every type. Hence $V_t \equiv V_t(\tau = 1)$ denotes the agent's willingness to pay for being in the treatment group. For simplicity we assume that there exists a known value $V_{\max} \in \mathbb{R} > 0$ such that for all $t \in T$, $V_t \in (-V_{\max}, V_{\max})$ and that the distribution over values induced by the distribution of types χ admits a density. The optimal effort for type t given treatment status τ is denoted by $e^*(\tau, t)$.²¹

4.1 Information Production in Open Selective Trials

A first benchmark result highlights the fact that selective trials are natural extensions of RCTs. An RCT is a mechanism $G_0 = (\emptyset, \pi_0)$. As $M = \emptyset$, no messages are sent, all agents are assigned to the treatment group with the same probability $\pi_0 \in (0, 1)$, and there are no transfers.

Fact 1 (full support sampling). *Consider a mechanism $G = (M, \mu)$. If there exists $\xi > 0$ such that for all $m \in M$, $\pi(m) \in (\xi, 1 - \xi)$, then $G_0 \preceq G$.*

Recalling that $\pi(m) \equiv \text{Prob}(\tau = 1|m)$, Fact 1 shows that if every type has a positive probability of being in the treatment or control group, then it is as informative as an RCT. Note this holds for any $\xi > 0$ because the sample size is infinite. We analyze small sample issues in Section 4.2.

As Plott and Zeiler (2005) and others have shown, information elicited in non-incentive compatible ways can be unreliable. Moreover, as Kremer and Miguel (2007) and others have noted, reported beliefs about a technology's return are often uncorrelated with use.

²¹At this stage, whether optimal effort is unique or not does not matter. We explicitly assume a unique optimal effort in Sections 5 and 6 to apply a convenient version of the Envelope Theorem.

Therefore, we focus on *strictly incentive compatible* assignment mechanisms—assignment mechanisms such that χ -almost every agent has a strictly preferred message.²²

Our next result shows that an open selective trial is a most informative trial if it identifies each agent’s value V_t , and, conditional on any expressed valuation, assigns a positive mass of agents to both the treatment and control group. Moreover, these are necessary conditions for an open selective trial to maximize informativeness.

Proposition 1 (most informative mechanisms). *Any strictly incentive compatible mechanism G identifies at most value V_t ($V_t = V_{t'} \Rightarrow m_G(t) = m_G(t')$).*

Whenever G identifies values V_t ($m_G(t) = m_G(t') \Rightarrow V_t = V_{t'}$) and satisfies full support ($0 < \inf_m \pi(m)$ and $\sup_m \pi(m) < 1$), then $G' \preceq G$ for any strictly incentive compatible mechanism G' .

It follows that open selective trials can at most identify the distribution of returns conditional on the agents’ valuation, which can be used to construct marginal treatment effects (MTEs). It is important to note that these mechanisms identify MTEs independently of the experimenter’s beliefs. The identification of MTEs in observational data requires a local instrument, that is, an instrument that changes the probability of adoption for agents with each possible value (Heckman and Vytlacil, 2005; Moffitt, 2008). Selective trials construct these local instruments by randomizing treatment conditional on an agent’s value. Hence, to the extent that elicited values are reliable, these mechanisms identify MTEs with a degree of robustness comparable to that with which RCTs identify average treatment effects.²³

²²Note that the mechanisms we consider can accommodate surveys. Consider the mechanism $G = (T, \pi_0)$ with message space $M = T$ where the likelihood of treatment is constant and equal to π_0 and no transfers are made. This is essentially an RCT supplemented with a rich survey. Note since assignment does not depend on the message, truthful revelation of one’s type is weakly dominant. Unfortunately, any other message is also weakly dominant. Hence, data generated by such a mechanism is likely to be unreliable, especially if figuring out one’s preferences is costly.

²³Note that selective trials also identify higher order moments of the outcome distribution conditional on treatment status and valuation, which may be useful to researchers.

Implementing Most Informative Trials. Here we exhibit two straightforward implementations of most informative selective trials.²⁴ The first is the BDM mechanisms described in Section 2.1, with the expanded message space $M = [-V_{\max}, V_{\max}]$. Once again the principal draws a price $p_i \in [-V_{\max}, V_{\max}]$ independently for each agent from a common c.d.f. F_p with support $[-V_{\max}, V_{\max}]$. If $m_i \geq p_i$, then the agent is assigned $(\tau = 1, p_i)$; otherwise, he is assigned $(\tau = 0, 0)$.

Fact 2 (BDM Implementation). *Whenever F_p has full support over $[-V_{\max}, V_{\max}]$, an agent with value V_t sends optimal message $m_{BDM} = V_t$ and the BDM mechanism is a most informative mechanism.*

A second implementation is a menu of lotteries. Consider mechanism G^* , where $M = (-\frac{1}{2}, \frac{1}{2})$, any agent sending message m is assigned to the treatment group with probability $\pi(m) = \frac{1}{2} + m$ and must make a transfer $p(m) = V_{\max}m^2$. One can think of agents as having a baseline probability of being in the treatment group equal to $\frac{1}{2}$ and deciding by how much they want to deviate from this baseline. An agent with value V_t chooses message m to maximize

$$\pi(m)V_t - p(m) = V_t \left(\frac{1}{2} + m \right) - V_{\max}m^2. \quad (3)$$

This problem is concave in m , and first order conditions yield an optimal message $V_t/2V_{\max}$ which identifies V_t . In addition, every agent is assigned to the treatment and control group with positive probability. Thus G^* is a most informative mechanism.

Note that G^* gives agents higher expected utility than an RCT which assigns agents to the treatment and control group with probability $\frac{1}{2}$. Indeed, for any RCT, a selective trial that assigns price $p = 0$ when π is the same as in the RCT will improve the expected utility of agents. Thus, selective trials may help decrease the number of agents who refuse randomization, which can approach 50% in medical trials (Jadad and Enkin, 2007).

²⁴See Appendix B for a description of selective trials that elicit coarser information using finite menus of lotteries.

4.2 The Cost of Running Selective Trials

In equilibrium, the menu of lotteries G^* yields sampling profile $\pi(V) = \frac{1}{2} \left(1 + \frac{V}{V_{\max}}\right)$, which is strictly increasing in value V . In the BDM mechanism the sampling profile, $\pi_{BDM}(V) = F_p(V)$, is also increasing in V . This is true of any mechanism.

Proposition 2 (monotonicity). *Consider a strictly incentive compatible mechanism G . If agents t and t' with $V_t > V_{t'}$ send messages $m_G(t) \neq m_G(t')$, then it must be that $\pi(m_G(t)) > \pi(m_G(t'))$.*

Thus, in any selective trial, agents with high values are over-sampled—they have a higher likelihood of being in the treatment group—and those with low values are under-sampled. In contrast, RCTs have a flat sampling profile. While sampling patterns do not matter when there is a large number of agents, they can significantly affect statistical power in small sample settings.

This issue is related to the recent development economics debate on charging for treatment in RCTs.²⁵ If, as in Ashraf et al. (forthcoming), willingness to pay is correlated with product usage, then eliciting willingness to pay might be quite useful in understanding true returns. If, instead, as in the case of Cohen and Dupas (2010), most agents have low values, and willingness to pay is a poor predictor of actual use, undersampling agents with low values may significantly reduce statistical power. Furthermore, in such a setting, willingness to pay provides little information about intended use.²⁶

We make two contributions to this debate. First, we note that when trade-offs between money and treatment are uninformative, selective trials can and should be based on more informative trade-offs. For instance, if most of the heterogeneity in willingness to pay is

²⁵Note that this literature is motivated by questions of efficiency, and is mostly interested in whether charging for usage improves how well treatment is matched with those who need and use it. This paper takes a slightly different perspective, and is interested in how controlling for willingness to pay improves inference from experimental trials.

²⁶As Dupas (2010) shows, this can also hinder social learning. Altogether, the Abdul Latif Jameel Poverty Action Lab recommends against charging prices for health technologies. For more details see <http://www.povertyactionlab.org/policy-lessons/health/pricing-health-products>.

driven by wealth and credit constraints, then eliciting willingness to wait, or willingness to perform a tedious task (like sitting through multiple information sessions) may be a better indicator of future usage than willingness to pay. If this is the case, selective trials can and should be designed around such tradeoffs. However, as we discuss in Section 7, this likely requires some knowledge of the agents and their environment. The technical details of extending our approach to non-monetary trade-offs can be found in Appendix A.

Second, we show that carefully designed selective trials can reduce the costs of oversampling by reducing the slope of the sampling profile.

Proposition 3 (sampling rates and incentives). *For any mechanism $G = (M, \mu)$ and $\underline{\rho} < \bar{\rho}$ in $(0, 1)$, there exists a mechanism $G' = (M, \mu')$ such that $G \preceq G'$, and for all $m \in M$, $\pi'(m) \in [\underline{\rho}, \bar{\rho}]$.*

The following must also hold. Denoting the expected utility of type t sending message m in mechanism G' (gross of transfers) by $U(t|m, G')$, then

$$\max_{m_1, m_2 \in M} |U(t|m_1, G') - U(t|m_2, G')| \leq 2(\bar{\rho} - \underline{\rho})V_{\max}.$$

Proposition 3 implies that it is always possible to reduce the slope of a mechanisms' sampling profile without affecting identification. Unfortunately, reducing the slope of the sampling profile also reduces incentives for truth-telling. We illustrate this with mechanisms $(G_\lambda^*)_{\lambda \in (0, 1)}$ which generalize G^* as follows: $M = (-\frac{1}{2}, \frac{1}{2})$, $\pi(m) = \frac{1}{2} + \lambda m$ and $p(m) = \lambda V_{\max} m^2$. As the slope of the sampling profile, λ , goes to zero, each agent will be sampled with probability approaching $\frac{1}{2}$ and will pay an amount approaching zero, irrespective of the message he sends. For any $\lambda > 0$, $m = V_t/2V_{\max}$ is still a dominant strategy for an agent of type t . However, if an agent with value V_t instead sends message $V/2V_{\max}$, his expected loss is

$$U(t|m = V_t/2V_{\max}) - U(t|m = V/2V_{\max}) = \frac{\lambda}{4V_{\max}}(V_t - V)^2,$$

which vanishes as the slope of the sampling profile λ goes to 0.

The important point is that despite this limitation, the slope of the sampling profile is a free parameter which the experimenter can and should optimize over. In particular, if the goal is merely to elicit willingness to pay rather than target the assignment of treatment, one can avoid excessive under-sampling of low value agents of the kind described by Cohen and Dupas (2010).

Altogether, this section has shown that open selective trials provide a simple way to identify MTEs and, more generally, the distribution of returns conditional on willingness to pay. In addition, while selective trials systematically oversample high value agents, this issue is negligible when sample size is large or agents are very responsive to incentives. However, as Section 2 highlighted, willingness to pay need not be a good predictor of actual effort and MTEs may not allow identification of deep parameters of interest. The following sections explore richer treatments which can better identify the role of effort.

5 Blind Selective Trials

5.1 Framework and Basic Results

In blind trials the agent is assigned a probability of being in the treatment group, $\phi \in [0, 1]$, which is disclosed to the agent, and an actual treatment status, $\tau \in \{0, 1\}$, which is known only to the principal. Thus, the pair (τ, ϕ) can be thought of as a full description of an agent's overall treatment. This class of selective blind trials nests both open trials (where $\phi \in \{0, 1\}$) and standard blind trials, where ϕ is fixed.

Assignment Mechanisms. As noted in Section 3, selective blind trials are mechanisms $G = (M, \mu)$ where $\mu : M \rightarrow \Delta([0, 1] \times \mathbb{R})$. Given a message m , μ assigns the agent a likelihood of being treated $\phi \in [0, 1]$ which is known to the agent, and a transfer $p \in \mathbb{R}$.

An actual treatment status $\tau \in \{0, 1\}$ is drawn according to ϕ . We denote by $\mu(\phi|m)$ the density over ϕ given message m .

Utility and Effort. An agent of type t 's value for uncertain treatment status ϕ is:

$$V_t(\phi) = \max_{e \in E} \left(\phi q_t(\tau=1, e) + (1-\phi)q_t(\tau=0, e) \right) \left(u(y=1, t) - u(y=0, t) \right) + u(y=0, t) - c(e, t). \quad (4)$$

The corresponding effort decision is $e^*(\phi, t)$, which we assume is unique.²⁷ Consistent with earlier notation we maintain $V_t(\phi=0) = 0$. Note that $V_t(\phi=1) = V_t$ is the agent's value for treatment in an open trial. Throughout the section, we keep ϕ as an argument of $V_t(\phi)$ and denote the value of $V_t(\phi)$ at φ by $V_t(\phi=\varphi)$. Thus, $V_t(\phi)$ denotes the entire mapping: $\varphi \mapsto V_t(\phi=\varphi)$.

Proposition 4 (most informative mechanisms). *Any strictly incentive compatible blind mechanism G identifies at most mapping $V_t(\phi)$ (that is, $V_t(\phi) = V_{t'}(\phi) \Rightarrow m_G(t) = m_G(t')$).*

If G identifies $V_t(\phi)$ (that is, $m_G(t) = m_G(t') \Rightarrow V_t(\phi) = V_{t'}(\phi)$) and satisfies $\inf_{\phi, m} \mu(\phi|m) > 0$ then $G' \preceq G$ for any strictly incentive compatible mechanism G' .

A simple generalization of the BDM mechanism is a most informative blind trial. Pick distributions, F_ϕ over $[0, 1]$, and $F_{p|\phi}$ over $[-V_{\max}, V_{\max}]$ with densities bounded away from 0. The blind BDM Mechanism (bBDM) has message space $M = [-V_{\max}, V_{\max}]^{[0,1]}$, so that a message m corresponds to a value function $V_t(\phi)$. Given message m , the principal draws values $\phi = \varphi$ and p according to distributions F_ϕ and $F_{p|\phi}$. If $m_i(\varphi) \geq p$, the agent is assigned (φ, p) . Otherwise, the agent is assigned $(0, 0)$. It is straightforward to show that $m_{bBDM}(t) = V_t(\phi)$. Additionally, bBDM satisfies the full sampling constraint $\inf_{\phi, m} \mu(\phi|m) > 0$.

Blind selective trials have two distinct advantages over open selective trials. First, blind selective trials decorrelate an agent's behavior and treatment status. As detailed in the next

²⁷Using the results of Milgrom and Segal (2002) this allows us to apply the usual Envelope Theorem to $V_t(\phi)$ in Proposition 6. Note that this also implies that $e^*(\phi, t)$ is continuous in ϕ .

subsection, this will allow the principal to identify whether empirical success rates are being driven by the agent’s behavior or by the treatment itself. Second, by identifying the value function $V_t(\phi)$, blind selective trials provide useful information about an agent’s intended behavior and his perceived success rate.

5.2 The Value of Decorrelating Beliefs and Treatment Status

Changes in success rates due to treatment come from two sources: the effect of the treatment itself, and the effect of behavioral changes associated with treatment. In an open trial, behavioral changes are perfectly correlated with changes in treatment status. As a result, the effect of treatment and the effect of behavioral changes induced by the expectation of treatment are hard to distinguish. In contrast, blind trials allow us to disentangle these two effects by distinguishing an agent’s actual treatment status τ and his (correct) belief ϕ that he is being treated.

To disentangle these effects, we focus on $\mathbb{E}[y|V_t(\phi), \phi = \varphi, \tau]$, the measured success rate conditional on the value function $V_t(\phi)$, belief $\phi = \varphi$ and treatment status τ , which is identified by selective blind trials. This allows identification of MTEs conditioned on the entire value function, $\Delta^{MTE}(V_t(\phi))$, as well as

$$\begin{aligned} \Delta^T(V_t(\phi)) &= \lim_{\substack{\varphi \rightarrow 0 \\ \varphi > 0}} \mathbb{E}[y|V_t(\phi), \phi = \varphi, \tau = 1] - \mathbb{E}[y|V_t(\phi), \phi = \varphi, \tau = 0] \\ \Delta^B(V_t(\phi)) &= \lim_{\substack{\varphi \rightarrow 1 \\ \varphi < 1}} \mathbb{E}[y|V_t(\phi), \phi = \varphi, \tau = 0] - \mathbb{E}[y|V_t(\phi), \phi = 0, \tau = 0]. \end{aligned}$$

As φ approaches zero, an agent’s effort converges to $e^*(\tau = 0, t)$, the effort he would expend if he knew he was not treated.²⁸ Hence, Δ^T identifies the returns to treatment keeping the agent’s behavior at its default level $e^*(\tau = 0, t)$. Similarly, as φ approaches one, the agent’s effort converges to $e^*(\tau = 1, t)$, the effort associated with sure treatment. Thus,

²⁸We must use a continuity argument because $\phi = 0$ implies $\tau = 0$, hence, there is no treatment group.

Δ^B is the effect of behavior change alone and

$$\Delta^I \equiv \Delta^{MTE} - \Delta^T - \Delta^B \tag{5}$$

measures the aggregate treatment effect (conditional on value $V_t(\phi)$), net of the effect of treatment and behavior alone. That is, Δ^I measures the interaction effect between behavior and treatment. If (5) is positive, then treatment and effort changes are complementary in producing successful outcomes. If, instead, (5) is negative, this suggests that there is a negative interaction between treatment and the *perceived* optimal effort of agents.²⁹

Being able to identify Δ^T and Δ^B has important practical implications. Consider, for example, a cholesterol-reducing drug. If subjects react to being in the study by eating more fatty foods, then the aggregate effect of treatment could be quite small even if the effect of the drug alone is significant. In this environment, Δ^T is the treatment effect purified of changes in behavior, that is, the effect of the drug on people who do not change their diet.

It is important to keep in mind when interpreting Δ^B and Δ^I that these are the direct and interaction effects at the agents' *perceived* optimal effort level $e^*(\tau = 1, t)$. In the example of the cholesterol reducing drug, the agent's perceived optimal behavior is to eat (rather than abstain from) fatty foods. Consequently, if the measured interaction Δ^I is small, this may be because effort does not improve the success rate of treatment, or because the agent is putting in low effort. In order to distinguish these two possibilities, we need additional information on the effort of agents. As the following subsection shows, this is what $V_t(\phi)$ provides.

²⁹These quantities can also be conditioned only on the value for sure treatment, V_t . Note also that Δ^T can be estimated using a standard blind RCT with a sufficiently low value of ϕ .

Note that selective blind trials can allow for double-blind designs in which the experimenter has varying beliefs over the likelihood that an agent is being treated. In some environments, varying the beliefs of the experimenter may help identify the treatment effect due to variations in experimenter behavior. However, treating this question properly requires a better understanding of the experimenter's incentive problem, which we abstract from in this paper.

5.3 The Value of Eliciting Preferences $V_t(\phi)$

As highlighted in Section 2.3, the mapping $V_t(\phi)$ can tell us whether and by how much treatment changes an agent's effort. Recalling that $V_t(\phi = 0) = 0$, knowledge of $V_t(\phi)$ provides the following simple test.

Proposition 5 (a test of “intention to change behavior”).

If $e^(\phi=0, t) = e^*(\phi=1, t)$, then for all φ , $V_t(\phi=\varphi) = \varphi V_t(\phi=1)$.*

If $e^(\phi=0, t) \neq e^*(\phi=1, t)$, then for all $\varphi \in (0, 1)$, $V_t(\phi=\varphi) < \varphi V_t(\phi=1)$.*

When effort changes with treatment status, the agent gets additional surplus from tailoring his behavior to τ . The difference $\varphi V_t(\phi=1) - V_t(\varphi)$ is thus the agent's willingness to pay to learn his actual treatment status, which will be zero if effort is independent of treatment. Recalling that $q_t(\tau, e)$ is an agent of type t 's perceived success rate, the value function $V_t(\phi)$ also allows us to estimate an agent's perceived returns to perceived optimal effort.

Proposition 6 (identifying perceived returns to effort). *For any value φ ,*

$$\left. \frac{\partial V_t(\phi)}{\partial \phi} \right|_{\varphi} = [q_t(\tau=1, e^*(\varphi, t)) - q_t(\tau=0, e^*(\varphi, t))] \times [u(y=1, t) - u(y=0, t)].$$

In particular, we can compute the ratio of perceived treatment effects, $q_t(\tau=1, e^*(\varphi, t)) - q_t(\tau=0, e^*(\varphi, t))$, at $\varphi = 1$ (with perceived optimal effort given treatment) and $\varphi = 0$ (with default effort):

$$\left. \frac{\partial V_t(\phi)}{\partial \phi} \right|_1 \bigg/ \left. \frac{\partial V_t(\phi)}{\partial \phi} \right|_0 = \frac{q_t(\tau=1, e^*(\varphi=1, t)) - q_t(\tau=0, e^*(\varphi=1, t))}{q_t(\tau=1, e^*(\varphi=0, t)) - q_t(\tau=0, e^*(\varphi=0, t))}. \quad (6)$$

This data helps us evaluate whether under-provision of effort is to blame for poor treatment effects. Returning to the example in Section 2, imagine a trial of a water treatment product known to the experimenter to be effective only if agents use it whenever they drink water. If measured returns to the treatment are low, there are two competing explanations:

1) the treatment is not effective in the agents' disease environment, 2) agents are not expending appropriate effort using the product. Agents' perceived returns can help distinguish these explanations. If perceived returns to effort are high, then the agent is likely to be expending significant effort, and it becomes more likely that the treatment is not effective in a particular disease environment. If, instead, perceived returns are low, it becomes more likely that the treatment has an effect that is unmeasured due to agents' lack of effort.

In addition, this data may provide some insight into the nature of placebo effects. If the indirect preferences of a group of agents indicates that they do not intend to change their behavior (for instance, via Proposition 5), yet exhibit positive behavioral effects ($\Delta^B > 0$), this indicates that the improvement due to a higher probability of uncertain treatment is affecting the results through unconscious, rather than conscious channels. If instead agents believe there are high returns to appropriate effort, this suggests that effect Δ^B is driven by conscious decisions.

6 Incentivized Selective Trials

We now show how quantities similar to those identified by blind selective trials can be identified without a placebo. This can be accomplished using an incentivized selective trial, which allows agents to express preferences over contracts.³⁰ A fully worked-out numerical example illustrating inference from incentivized trials is given in Appendix C.

6.1 Framework and Basic Results

Assignment Mechanisms. As noted in Section 3, an incentivized trial is a mechanism $G = (M, \mu)$, where $\mu : M \rightarrow \Delta(\{0, 1\} \times \mathbb{R} \times \mathbb{R})$. Given a message m , μ is used to draw a treatment status τ , a fixed transfer p from the agent, as well as a bonus w (both of which

³⁰For field experiments using explicit incentives see, for instance, Gertler (2004); Schultz (2004); Volpp et al. (2006, 2008); Thornton (2008); Kremer et al. (2009).

may be negative in the case of insurance) transferred to the agent in the event of success. The pair (τ, w) can be thought of as an aggregate treatment.

Utility and Effort. The agents' indirect preferences over contracts (τ, w) , denoted by $V_t(\tau, w)$, are given by

$$V_t(\tau, w) = \max_{e \in E} q_t(\tau, e)[u(y=1, t) + w] + [1 - q_t(\tau, e)]u(y=0, t) - c(e, t). \quad (7)$$

We denote by $e^*(\tau, w, t)$ the induced effort level, and maintain the normalization $V_t(\tau = 0, w = 0) = 0$.

Insurance. A specific value w that will be useful is $w_0 \equiv -[u(y = 1, t) - u(y = 0, t)]$. When the agent receives a positive baseline transfer for participating in the experiment, the negative bonus w_0 essentially provides the agent with perfect insurance over the outcome y . When fully insured, the agent will put in the effort that minimizes the cost of his effort regardless of his treatment status. Note that this level of effort differs from the default behavior of untreated agents in an open trial, as agents in open trials may still be exerting some effort to improve their outcomes.

We proceed by assuming that w_0 is known to the principal. At the end of the section we show w_0 can be inferred from the mapping $V_t(\tau, w)$ under fairly mild conditions. Alternatively, as w_0 is the monetary value of success, it could be calibrated from other data.

6.2 What can be Inferred from Incentivized Trials?

It is straightforward to extend Propositions 1 and 4, which characterize most informative mechanisms. That is, G is a most informative incentivized trial if it identifies $V_t(\tau, w)$ and, given any message, puts positive density on all possible treatments (τ, w) . As before, the *BDM* mechanism can be adapted to identify $V_t(\tau, w)$. Note that the information produced

by incentivized trials nests that produced by open trials. In particular, $V_t(\tau=1, w=0) = V_t$.

As in the case of blind selective trials, incentivized selective trials allow us to decorrelate treatment and effort, as well as infer an agent's perceptions of how effort affects outcomes. Incentivized selective trials recover the empirical success rate $\mathbb{E}[y|V(\tau, w), \tau, w]$ as a function of preferences, treatment and incentives. This will be independent of reward w if effort does not matter for outcomes or if incentives do not affect effort provision.

Isolating returns to treatment and returns to effort. A contract with transfer $w_0 \equiv -[u(y=1, t) - u(y=0, t)]$ provides the agent with perfect insurance. The optimal effort given full insurance will be the same regardless of treatment status. Given w_0 , we can identify two quantities similar to those discussed in Section 5.2:

$$\begin{aligned} \text{Returns to Treatment | No Effort} &= \mathbb{E}[y|V_t(\tau, w), \tau=1, w=w_0] - \mathbb{E}[y|V_t(\tau, w), \tau=0, w=w_0] \\ \text{Returns to Effort | Treatment} &= \mathbb{E}[y|V_t(\tau, w), \tau=1, w=0] - \mathbb{E}[y|V_t(\tau, w), \tau=1, w=w_0] \end{aligned}$$

Note that here, returns are measured expending minimal effort $e^*(\tau, w_0, t)$ as a baseline, rather than the default effort level $e^*(\tau=0, w=0, t)$ exerted by agents in the control group of an open trial.

Identifying Perceived Returns to Effort. Indirect preferences $V_t(\tau, w)$ also provide a handle on perceived returns to effort. Recall that $q_t(\tau, e)$ denotes the agent's perceived likelihood of success given treatment status τ and effort e .

Proposition 7 (identifying perceived success rates).

$$\forall \tau, w, \quad \frac{\partial V_t(\tau, w)}{\partial w} = q_t(\tau, e^*(\tau, w, t)).$$

Given knowledge of w_0 , this allows us to compute subjective returns to treatment and per-

ceived appropriate effort:

$$\text{Perceived Returns to Treatment} = q_t(\tau=1, w=w_0|V_t(\tau, w)) - q_t(\tau=0, w=w_0|V_t(\tau, w))$$

$$\text{Perceived Returns to Effort} = q_t(\tau=1, w=0|V_t(\tau, w)) - q_t(\tau=1, w=w_0|V_t(\tau, w)).$$

Note that if perceived returns to effort are low, this can indicate that an agent plans on putting little or no effort into using the technology. The principal can use this information in deciding which agents' usage to monitor more closely.

The monetary equivalent of the cost of effort agents incur to obtain the perceived return to effort above can be obtained by rearranging (7):

$$c(e^*(\tau, w=0, t)) - c(e^*(\tau, w=w_0, t)) = -w_0 \times q_t(\tau, e^*(\tau, w=0, t)) - [V_t(\tau, w=0) - V_t(\tau, w=w_0)].$$

Note that all parameters on the right hand side are identified from data, except perhaps w_0 . Identifying the costs incurred by agents can greatly improve inference. In particular it allows us to distinguish—among agents who think that appropriate effort has high returns—those who think that only a small amount of effort is sufficient to obtain high returns, from those who think that a significant amount of effort is necessary to obtain high returns.

Identifying the full insurance contract. One drawback of incentivized trials is that they rely on identifying the full insurance contract w_0 . However, this quantity can be identified under mild additional assumptions.

Fact 3. *If outcome $y = 1$ yields strictly greater utility than $y = 0$: $u(y = 1, t) > u(y = 0, t)$ and agents perceive treatment to always be beneficial:*

$$\forall e_0 \in E, \exists e_1 \in E \text{ s.t. } c(e_1, t) \leq c(e_0, t) \quad \text{and} \quad q_t(\tau=0, e_0) < q_t(\tau=1, e_1)$$

$$\text{then } w_0 = \min\{w \mid \forall w' > w, V_t(\tau=1, w') > V_t(\tau=0, w')\}.$$

In words, when treatment facilitates success, the full insurance transfer w_0 is the highest transfer such that the agent no longer values obtaining the treatment. Note that our assumptions rule out cases where the agent believes he may be hurt by the treatment, as well as environments where the agent only values treatment for reasons other than its impact on the experimenter’s outcome of interest. Whenever the assumptions do not hold, w_0 must be calibrated from alternative data (for example, the expected amount of wages lost when sick).

7 Discussion

This paper studies the inference and external validity concerns that arise when experimental subjects take unobserved decisions which can affect outcomes. In particular, since effort provision is driven by beliefs and beliefs can respond to information, the returns measured by an RCT may not be representative of the returns that a better informed population would obtain. To address this issue we take a principal-agent approach to trial design where the principal maximizes the informativeness of data. This leads us to study selective trials, which improve on RCTs by letting agents express preferences over treatments of varying richness. We show that selective trials can identify whether agents’ beliefs are reducing measured treatment effects and separate the returns from treatment, from effort, and from their interaction.

More generally, this paper advocates a mechanism design approach to randomized controlled trials, which we believe can help build bridges between purely frequentist methods—largely concerned with robustness and internal validity—and structural methods—which use models to identify deep parameters necessary to evaluate external validity. While we believe that this research agenda can yield many useful applications, successfully implementing its insights is likely to pose a number of interesting challenges, some of which we delineate below.

Implementation Issues. In theory, the selective trials described in this paper are robust and require no specific knowledge on the part of the principal.³¹ However, there are many issues to consider when implementing them.

The first is that trade-offs with respect to money need not be the most informative, and appropriate local knowledge is needed to pick trade-offs that are informative of the agents' intended behavior. For instance, as we highlight in Section 4, when agents are credit constrained, monetary trade-offs may be less informative than willingness to wait or willingness to perform tedious tasks.

A second, more pernicious, difficulty is that boundedly rational agents might not play dominant strategies, so that the messages they send are not good signals of their actual values (Keller et al., 1993; Bohm et al., 1997). This issue can be mitigated by giving agents multiple opportunities to learn how to play the mechanism used to elicit their valuation before they are asked to express preferences over treatment (Plott and Zeiler, 2005). In addition, as Section 4 highlights, there are many possible implementations of selective trials some of which may be more appropriate than others in the field.

Another delicate implementation issue is that if the act of making choices changes agents' preferences, selective trials may introduce additional noise. For instance, imagine that when agents express a strong desire for treatment but do not get it, they will attempt to obtain treatment by other means, but would not do so if valuation were never elicited. In this setting, running a selective trial may prevent the experimenter from building an appropriate control group.³²

Finally, the fact that agents may make inferences from the principal's choice of experimental design may interfere with their behavior. For instance, similar to Milgrom and Roberts (1986a), if treatment is only available at a high cost, agents may infer that the tech-

³¹Even the requirement that values be bounded above and below can be dispensed in theory in BDM mechanisms by picking a distribution of prices with an unbounded support.

³²Of course, in a standard open RCT, agents in the control group who highly value the treatment may try to obtain it. A selective trial may decrease noise by identifying these agents and facilitating their monitoring.

nology is more valuable. In such environments, a careful principal should take into account how experimental design influences behavior before drawing inferences. In addition, in the spirit of Milgrom and Roberts (1986b); Mullainathan et al. (2008); Rayo and Segal (2008), or Kamenica and Gentzkow (2009) a principal may want to optimize the structure of her trial, as well as the information it generates, to improve future adoption by agents.

Extension to Dynamic Mechanisms. Throughout the paper, we focus on mechanisms which allow us to elicit agents' preferences only once. Note that this could occur either before or after an agent has been exposed to the technology. In environments where agents learn over the course of the experiment, it may be valuable to elicit the agents' preferences over time.³³ Consider a technology that requires sustained effort to yield returns, for example: anti-depressants with delayed effects, technologies exhibiting significant learning-by-doing, and so on. Eliciting preferences over time may improve inference by helping to distinguish agents exhibiting consistent motivation throughout the trial from agents whose motivation drops in the middle. However, eliciting preferences over time may be complicated when anticipated treatment status changes current effort expenditure. In particular, if an agent is promised treatment in future periods to induce a particular current effort level, then it becomes impossible to elicit preferences in the future without breaking this promise. We leave the analysis of such mechanisms and their limitations for future work.

Extension to multi-agent mechanisms. The mechanisms considered in this paper are all single-agent mechanisms—an agent's assignment depends only on the message he sends and not on the messages sent by others. This allows us to identify the agent's preferences as well as his beliefs over his own treatment effects and returns to effort. Considering multi-agent mechanisms, in which assignment depends on the messages sent by others, can allow

³³Philipson and Desimone (1997); Philipson and Hedges (1998); Scharfstein et al. (1999); Chan and Hamilton (2006) take a different approach to this problem by incorporating the information from observable non-compliance into econometric models in order to estimate more accurate treatment effects.

us to identify the agent's beliefs about others' values, others' success rates, and so on.

The information elicited by multiple-agent mechanisms may be useful if there are externalities between agents, as in Miguel and Kremer (2004), or to get a tighter handle on social learning. For example, if we observe that most agents have low value for the technology but believe that others have high value for the technology, this suggests a specific failure of social learning, and provides us with the means to correct it. Indeed, if most agents do not expend effort using the technology but believe others do, then agents will interpret others' poor outcomes as a signal that even with high effort the technology does not yield returns. Providing the agents with actual data on others' willingness to pay corrects these inference mistakes and may increase experimentation. Again, we leave more in depth analysis of such mechanisms to future work.

Appendix — Not for Publication

A Extensions

A.1 General Outcome Space

Most of the results extend directly to the case where y takes values in a general outcome space Y , and is distributed according to some density function $f_y(R, \tau, e, t)$. We denote by $f_{y,t}(\tau, e) \equiv \int_R f_y(R, \tau, e, t) dt(R)$ the subjective distribution of returns from the perspective of an agent of type t . Values simply go from being sums of two terms to being integrals. The only change to the mechanisms we consider is that incentive contracts are now functions $w : Y \rightarrow \mathbb{R}$. Indeed, we have that

$$\begin{aligned} V_t &= \max_{e \in E} \int_y u(y, t) f_{y,t}(\tau = 1, e) dy - c(e, t) \\ V_t(\phi) &= \max_{e \in E} \phi \int_y u(y, t) f_{y,t}(\tau = 1, e) dy + (1 - \phi) \int_y u(y, t) f_{y,t}(\tau = 0, e) dy - c(e, t) \\ V_t(\tau, w) &= \max_{e \in E} \int_y [u(y, t) + w(y)] f_{y,t}(\tau, e) dy - c(e, t). \end{aligned}$$

Propositions 1, 2, 3, 4 and 5 extend directly with these generalized value functions. Propositions 6 and 7, which identify subjective returns differ as follows. Proposition 7, which deals with incentivized trials is the easiest to extend. Indeed, we have that

$$\forall y_0, \quad \frac{\partial V_t(\tau, w)}{\partial w(y_0)} = f_{y,t}(\tau, e^*(\tau, w, t))(y_0),$$

which is a direct extension of Proposition 7.

Proposition 6, which deals with blind trials is more difficult to extend as now we have only a one-dimensional instrument, $\phi \in [0, 1]$ to identify an entire function $f_{y,t}$ rather than

the single parameter q_t . We now have that

$$\frac{\partial V_t(\phi)}{\partial \phi} = \int_y u(y, t) [f_{y,t}(\tau = 1, e^*(\phi, t))(y) - f_{y,t}(\tau = 0, e^*(\phi, t))(y)] dy \quad (8)$$

which corresponds to a utility weighted subjective treatment effect given subjectively appropriate effort under belief ϕ .

A.2 Eliciting Preferences under Non-Quasilinear Utility

The approach developed in this paper largely extends to the case where preferences are not quasilinear, although we must consider slightly different mechanisms. We now consider utility taking the form $u(y, e, p, t)$ where $y \in Y$, $e \in E$, p is a monetary outcome and t is the agent's type.

In the case of open trials, indirect preferences take the following form:

$$V_t(\tau, p) = \max_e \int_y u(y, e, p, t) f_{y,t}(\tau, e) dy.$$

Say we want to elicit preference over a range $(\tau, p) \in \{0, 1\} \times [\underline{p}, \bar{p}]$. We assume for simplicity that for all such (τ, p) , $V_t(\tau = 0, \bar{p}) \leq V_t(\tau, p) < V_t(\tau = 1, \underline{p})$. We normalize $V_t(\tau = 0, p = \bar{p}) = 0$ and $V_t(\tau = 1, p = \underline{p}) = 1$. Consider the following generalization of BDM mechanism: the agent sends a message $m \in \mathbb{R}^{\{0,1\} \times [\underline{p}, \bar{p}]}$, which corresponds to a value function; the principal randomly picks (τ, p, λ) from some continuous distribution over $\{0, 1\} \times [\underline{p}, \bar{p}] \times [0, 1]$; the agent is assigned (τ, p) if $m(\tau, p) > \lambda$ and the lottery $\lambda \times (\tau = 1, p = \underline{p}) + (1 - \lambda) \times (\tau = 0, p = \bar{p})$ otherwise. In this setting it is dominant for the agent to send message $m = V_t$. Similar mechanisms allow us to identify indirect preferences in the case of blind and incentivized trials.

Propositions 1, 2, 3, 4 and 5 extend directly with these generalized value functions. Again, extending Propositions 6 and 7 requires some more work. Proposition 6—which identifies

subjective returns to effort using blind trials—extends as is when $y \in \{0, 1\}$, and extends according to equation (8) when y takes values in a general outcome set Y . Proposition 7 extends as is when preferences are separable in money, that is, when $u(y, e, p, t) = u_0(y, e, t) - u_1(p, t)$. When preferences are not separable in money, incentivized trials allow us to identify $f_{y,t}(y) \frac{\partial u}{\partial p|_{y,p}}$ for all values of y and p (when preferences are separable, the multiplicative constant can be identified from the fact that probabilities sum to 1).

B Implementation of Open Selective Trials as Finite Menus of Lotteries.

The mechanisms described in the paper all use a continuum of messages and elicit the agent's exact willingness to pay. Of course it is possible to elicit coarser information. This example shows how to identify which of N intervals an agent's willingness to pay belongs to.

Let the experimenter pick value thresholds $-V_{\max} = V_0 < V_1 < \dots < V_N = V_{\max}$. She can elicit the interval where an agent's value lies by offering a menu of lotteries. This menu is constructed with messages $M = \{1, \dots, N\}$ and any increasing sequence $\pi(1) < \pi(2) < \dots < \pi(N)$ of sampling rates. Thus, message $m \in M$ corresponds to buying the lottery that delivers the treatment with probability $\pi(m)$. In order to match these messages with the appropriate value interval, the experimenter simply sets $p(m)$, the price of lottery m , according to:

$$\forall k > 1, \quad p(k) = p(k-1) + (\pi(k) - \pi(k-1))V_{k-1}. \quad (9)$$

Note that the sequence of prices is entirely determined by $p(1)$. Denote by $G^{\pi,p}$ the mechanism corresponding to this menu of lotteries, then:

Fact 4. $m_{G^{\pi,p}}(V_t) = k$ if and only if $V_t \in [V_{k-1}, V_k]$.

This emphasizes the many degrees of freedom the experimenter has when implementing selective trials as menus of lotteries. The value intervals according to which agents are

classified, and the rates according to which they obtain treatment are, to a large extent, free parameters. Still, sampling rates are increasing in an agent’s value.

C A Numerical Example Illustrating Inference from Incentivized Trials

This section illustrates step by step the process of inference from trial data, starting with a standard RCT, adding data from open selective trials and concluding by adding both objective and subjective data from an incentivized trial.

As regards the environment, we return to a setting where returns are two dimensional: $R = (R_b, R_e)$. As before, in the context of a water treatment product, R_b could be the baseline returns of using the water treatment product only when it is convenient to do so and R_e the returns to using it more thoroughly (for instance, bringing treated water when away from home). Success rates are given by:

$$q(\tau=0, e) = 0 \quad \text{and} \quad q(\tau=1, e) = R_b + eR_e,$$

where $e \in \mathbb{R}_+$ is the agent’s effort provision. An agent with type t has beliefs $R_t = (R_{b,t}, R_{e,t})$ and maximizes $\mathbb{E}_t[y] - c(e)$ where $c(e) = \frac{e^2}{2}$. The effort expended in an incentivized trial is thus $e^*(w, t) = R_{e,t}(1 + w)$, which nests the effort decision of an open trial, $e^*(w=0, t) = R_{e,t}$.

Throughout, we illustrate the inference process by considering the case where each parameter has a low and high value: $R_e, R_{e,t} \in \{1/4, 1/2\}$, $R_b \in \{0, 1/8\}$ and $R_{b,t} \in \{0, 3/32\}$. Each element of a selective trial adds data which will narrow the set of possible values.³⁴

³⁴For simplicity, we consider priors that put point masses on few possible states. Unfortunately, such strong priors can result in degenerate inference problems. We computed the states to keep the inference problem well defined and better reflect the mechanics of inference from a continuous state space. This accounts for the somewhat unusual aspect of our parameter values.

Inference from an RCT. An RCT identifies the average treatment effect, $\widehat{\Delta} = R_b + R_e \times R_{e,t}$. For the numerical values specified above the possible outcomes are described in the following matrix

	$R_e = 1/2$		$R_e = 1/4$	
	$R_{e,t} = 1/2$	$R_{e,t} = 1/4$	$R_{e,t} = 1/2$	$R_{e,t} = 1/4$
$R_b = 1/8$	$\widehat{\Delta} = 3/8$	$\widehat{\Delta} = 1/4$	$\widehat{\Delta} = 1/4$	$\widehat{\Delta} = 3/16$
$R_b = 0$	$\widehat{\Delta} = 1/4$	$\widehat{\Delta} = 1/8$	$\widehat{\Delta} = 1/8$	$\widehat{\Delta} = 1/16$

As illustrated by the matrix, if $\widehat{\Delta} \in \{1/16, 3/16, 3/8\}$ this identifies the returns of the technology (R_b, R_e) . However, treatment effects $\widehat{\Delta} \in \{1/8, 1/4\}$ are consistent with multiple true returns.³⁵ In particular, when $\widehat{\Delta} = 1/4$, it may be that casual use of the water treatment product is not particularly effective ($R_b = 0$), more thorough use is not particularly effective ($R_e = 1/4$), or more thorough use is effective, but agents don't believe it is, and so do not put much effort into using the water treatment product more thoroughly ($R_e = 1/4, R_{e,t} = 1/2$).

Inference from a Selective Open Trial. By Fact 1, open selective trials identify treatment effects $\widehat{\Delta}$. Additionally, by Proposition 1, an open selective trial identifies the agent's willingness to pay for treatment $V_t = R_{b,t} + R_{e,t}^2/2$. To illustrate the value of this data, focus on the case where $\widehat{\Delta} = 1/4$. As shown above, this is consistent with three different vectors of $(R_b, R_e, \text{ and } R_{e,t})$. Based on this, we illustrate the six possible values of V_t in the following matrix:

³⁵For example, $(R_b = 0, R_e = 1/2, R_{e,t} = 1/2), (R_b = 1/8, R_e = 1/2, R_{e,t} = 1/4)$ and $(R_b = 1/8, R_e = 1/4, R_{e,t} = 1/2)$ are all consistent with $\widehat{\Delta} = 1/4$.

Note that agents' beliefs may be self confirming. For instance, an agent who believes that effort has high returns, $R_{e,t} = 1/2$, who observes $\widehat{\Delta} = 1/4$ will continue to believe returns are high, even though this data could be generated by $R_e = 1/4$. Such self-confirming beliefs are frequent in the experimentation and social learning literatures. See for instance Rothschild (1974); Banerjee (1992); Bikhchandani et al. (1992).

	$R_b = 0, R_e = 1/2, R_{e,t} = 1/2$	$R_b = 1/8, R_e = 1/2, R_{e,t} = 1/4$	$R_b = 1/8, R_e = 1/4, R_{e,t} = 1/2$
$R_{b,t} = 3/32$	$V_t = 7/32$	$V_t = 1/8$	$V_t = 7/32$
$R_{b,t} = 0$	$V_t = 1/8$	$V_t = 1/32$	$V_t = 1/8$

If $V_t = 1/32$ the data from selective trials indicates $R_{e,t} = 1/4 = e^*$. As the treatment effect is $\widehat{\Delta} = 1/4$ the only consistent returns are $R_b = 1/8$ and $R_e = 1/2$. If $V_t = 7/32$, there remains uncertainty since the data is consistent with both $(R_b = 0, R_e = 1/2)$ and $(R_b = 0, R_e = 1/4)$. Finally if $V_t = 1/8$, the data is consistent with any of the states $(R_b, R_e, R_{e,t})$ that produce $\widehat{\Delta} = 1/4$. That is to say that even in this limited example, data from a selective open trial (and hence, MTEs) may not help in identifying underlying returns. We now turn to how incentivized trials allow us to infer whether effort or returns to effort are low.

Inference from an Incentivized Trial. Incentivized trials yield:

$$\widehat{\Delta}(w) = R_b + R_e \times R_{e,t}(1+w) \quad \text{and} \quad V_t(\tau=1, w) = R_{b,t}(1+w) + \frac{[R_{e,t}(1+w)]^2}{2}.$$

As an open selective trial already identifies $V_t = V_t(w=0) = R_{b,t} + R_{e,t}^2/2$ and $\widehat{\Delta} = \widehat{\Delta}(w=0) = R_b + R_e \times R_{e,t}$, by eliciting valuations and treatment effects for a small w the principal can also identify $\left. \frac{\partial V_t(\tau, w)}{\partial w} \right|_{w=0} = R_{b,t} + R_{e,t}$ and $\left. \frac{\partial \widehat{\Delta}(w)}{\partial w} \right|_{w=0} = R_e \times R_{e,t}$. With this data the principal can identify:

$$R_{e,t} = 1 - \left[1 + 2 \left(\left. \frac{\partial V_t}{\partial w} \right|_{w=0} - V_t(w=0) \right) \right]^{1/2}$$

and thus, the rest of the unknown parameters: $R_e = \left. \frac{\partial \widehat{\Delta}(w)}{\partial w} \right|_{w=0} / R_{e,t}$, $R_{b,t} = \left. \frac{\partial V_t(\tau, w)}{\partial w} \right|_{w=0} - R_{e,t}$, $R_b = \widehat{\Delta} - R_e \times R_{e,t}$. The same information can be identified in a mathematically simpler, but more data intensive, way by identifying w_0 and the empirical quantities associated with that value.

Altogether, incentivized selective trials allow us to identify both the true returns (R_b, R_e) and the agents' beliefs $(R_{b,t}, R_{e,t})$. Thus, in this example, data from a selective incentivized trial allows an experimenter to determine how effective casual and thorough use of the water treatment product is without having to observe individual agents' usage. This is possible as eliciting each agents' indirect preferences over the water treatment product and bonuses associated with staying healthy allows the experimenter to infer the agents' beliefs about the effects of casual and more thorough usage. This, in turn, allows the experimenter to infer behavior and identify the deep structural parameters determining the product's effectiveness, as well as how beliefs about effectiveness lead to different outcomes.

D Proofs

Proof of Fact 1: We break down the data \mathbf{d}_G in two subsamples $(d_G^{\sigma_0(i)})_{i \in \mathbb{N}}$ and $(d_G^{\sigma_1(i)})_{i \in \mathbb{N}}$ such that σ_0, σ_1 are non-decreasing mappings from \mathbb{N} to \mathbb{N} , and for all $i \in \mathbb{N}$, $\tau_{\sigma_0(i)} = 0$ and $\tau_{\sigma_1(i)} = 1$. Since $\forall m, \pi(m) \in [\xi, 1 - \xi]$, we have that each such subsample is infinite and we can pick σ_1 and σ_0 to be strictly increasing from \mathbb{N} to \mathbb{N} . We define mapping h (such that $h(\mathbf{d}_G) \sim \mathbf{d}_{G_0}$) as follows.

We use the notation $h(\mathbf{d}_G) = (d_i^h)_{i \in \mathbb{N}}$, where $d_i^h = (m_i^h, p_i^h, \tau_i^h, y_i^h)$. For every $i \in \mathbb{N}$, we set $m_i^h = \emptyset$, $p_i^h = 0$, we draw τ_i^h as the Bernoulli variable of parameter π_0 . Finally we set $y_i^h = y_{\sigma_{\tau_i^h}(i)}$. It is easy to check that indeed, $h(\mathbf{d}_G) \sim \mathbf{d}_{G_0}$. ■

Proof of Proposition 1: The proof of the first claim is very similar to that of Fact 1. Consider a mechanism $G = (M, \mu_G)$ such that every player has a strictly dominant strategy. An agent with value $V(t_i)$ chooses a message m_i to solve

$$\max_{m \in M} \pi(m)V(t_i) - \mathbb{E}_\mu(p_i | m_i = m).$$

This problem is entirely defined by player i 's value $V(t_i)$. Since a.e. player has a strictly optimal message, this problem has a unique solution for a.e. value.

We now construct a mapping $h : \mathcal{D} \rightarrow \Delta(\mathcal{D})$ such that the data generated by G' can be simulated from data generated by G using mapping h . For simplicity we describe the mapping h in the case where M is finite. Given \mathbf{d}_G , $h(\mathbf{d}_G)$ is generated as follows.

First, we break down the basic data \mathbf{d}_G in $2 \times \text{card } M$ subsets, according to treatment τ and the message $m_G(V)$ corresponding to the value declared by the agent. Formally, for all $m \in M$ and $\tau \in \{0, 1\}$, we define $(d_G^{\sigma_{m,\tau}(i)})_{i \in \mathbb{N}}$ the ordered subsequence such that for all i , $m_G(V_{\sigma_{m,\tau}(i)}) = m$ and $\tau_{\sigma_{m,\tau}(i)} = \tau$. Since $0 < \inf_m \pi(m) < \sup_m \pi(m) < 1$, all these subsamples are infinite. Hence, $\sigma_{m,\tau}$ can be chosen to be strictly increasing from $\mathbb{N} \rightarrow \mathbb{N}$. We use these subsamples to simulate data $\mathbf{d}_{G'}$.

Let us denote $h(\mathbf{d}_G) = (d_i^h)_{i \in \mathbb{N}}$. For all $i \in \mathbb{N}$, $d_i^h = (m_i^h, p_i^h, \tau_i^h, y_i^h)$. We first set $m_i^h = m_{G'}(V_i)$. Then using $\mu_{G'}(m_i^h)$, we draw values τ_i^h and p_i^h . Finally we set $y_i^h = y_{\sigma_{m_i^h, \tau_i^h}(i)}$. This defines $h : \mathcal{D} \rightarrow \Delta(\mathcal{D})$. It is easy to check that $h(\mathbf{d}_G) \sim \mathbf{d}_{G'}$.³⁶ This concludes the proof. ■

Proof of Fact 2: The fact that the BDM mechanism elicits values is well-known. Since F_p has full support over $[-V_{\max}, V_{\max}]$, assignment to treatment also satisfies full-support and the second part of Proposition 1 implies that G_{BDM} is a most informative mechanism. ■

Proof of Proposition 2: Agents of type t and t' are such that $V_t > V_{t'}$ and $m_G(t) \neq m_G(t')$. Denote $\pi(m) = \mu_G(\tau = 1|m)$ and $p_m = \mathbb{E}_{\mu_G(\cdot|m)} p$. It must be that

$$\begin{aligned} \pi(m_G(t))V_t - p_{m_G(t)} &> \pi(m_G(t))V_{t'} - p_{m_G(t)} \\ \pi(m_G(t'))V_{t'} - p_{m_G(t')} &> \pi(m_G(t'))V_t - p_{m_G(t')}. \end{aligned}$$

³⁶Note that for the sake of notational simplicity, this construction ends up wasting data points by not taking consecutive elements from the subsamples. This is inconsequential here since we have infinitely many data points.

Adding the two inequalities yields that $[\pi(m_G(t)) - \pi(m_G(t'))](V_t - V_{t'}) > 0$, which implies that $\pi(m_G(t)) > \pi(m_G(t'))$. ■

Proof of Proposition 3: We begin with the first assertion. Given mechanism $G = (M, \mu)$, we define mechanism $G' = (M, \mu')$ as follows:

$$\forall m \in M, \quad \mu'(m) = \begin{cases} \tau = 0, p = 0 & \text{with probability } \underline{\rho} \\ \mu(m) & \text{with probability } \bar{\rho} - \underline{\rho} \\ \tau = 1, p = 0 & \text{with probability } \bar{\rho} \end{cases}$$

Clearly mechanism G' is strategically equivalent to mechanism G . The proof that $G \preceq G'$ is omitted since it is essentially identical to that of Fact 1.

We now turn to the second assertion. Consider two messages m_1 and m_2 respectively (and optimally) sent by types with values V_1 and V_2 . Let $p_m = \mathbb{E}_{\mu_G(\cdot|m)}p$. We must have that

$$\begin{aligned} \pi_{G'}(m_1)V_1 - p_{G'}(m_1) &\geq \pi_{G'}m_2V_1 - p_{G'}(m_2) \\ \pi_{G'}(m_2)V_2 - p_{G'}(m_2) &\geq \pi_{G'}(m_1)V_2 - p_{G'}(m_1). \end{aligned}$$

These two inequalities yield that $(\pi_{G'}m_2 - \pi_{G'}m_1)V_1 \leq p_{G'}(m_2) - p_{G'}(m_1) \leq (\pi_{G'}(m_2) - \pi_{G'}(m_1))V_2$, which implies that $|p_{G'}(m_2) - p_{G'}(m_1)| < (\bar{\rho} - \underline{\rho})V_{\max}$. Hence the difference in utilities between sending two messages m_1 and m_2 for an agent with value $V \in [-V_{\max}, V_{\max}]$ is $|(\pi_{G'}(m_1) - \pi_{G'}(m_2))V - p_{G'}(m_1) + p_{G'}(m_2)| \leq 2(\bar{\rho} - \underline{\rho})V_{\max}$. ■

Proof of Proposition 4: The proof of Proposition 4 is essentially identical to that of Proposition 1 and hence omitted. ■

Proof of Proposition 5: The proof is given for the general case where there might be mul-

multiple optimal effort choices. Let $V_t(\tau, e)$ denote the expected value of type t under treatment status τ and when putting effort e . We have that

$$\begin{aligned} V_t(\phi) &= \max_{e \in E} \phi V_t(\tau=1, e) + (1 - \phi) V_t(\tau=0, e) \\ &\leq \phi \max_{e \in E} V_t(\tau=1, e) + (1 - \phi) \max_{e \in E} V_t(\tau=0, e). \end{aligned}$$

If $\arg \max_{e \in E} V_t(\tau=1, e) \cap \arg \max_{e \in E} V_t(\tau=0, e) \neq \emptyset$, the inequality is an equality and, since we normalized $V_t(\phi=0) = 0$ we obtain that $V_t(\phi) = \phi V_t(\phi=1)$. Inversely, if $\arg \max_{e \in E} V_t(\tau=1, e) \cap \arg \max_{e \in E} V_t(\tau=0, e) = \emptyset$, the inequality is strict and $V_t(\phi) < \phi V_t(\phi=1)$. ■

Proof of Proposition 6: The result follows directly from applying the Envelope Theorem to equation (4). ■

Proof of Proposition 7: The result follows directly from applying the Envelope Theorem to equation (7). ■

Proof of Fact 3: Whenever $w = w_0$, the agent is perfectly insured and $V_t(\tau=1, w) = V_t(\tau=0, w)$ since access to the technology is valuable only in so far as it affects outcomes. We now show that whenever $w > w_0$, $V_t(\tau=1, w) > V_t(\tau=0, w)$. The agent's value is

$$V_t(\tau, w) = \max_{e \in E} q_t(\tau, e) [u(y=1, t) - u(y=0, t) + w] + u(y=0, t) - c(e, t).$$

Let e_0^* be the agent's optimal effort level if $\tau = 0$. By assumption, there exists e_1 such that $c(e_1, t) \leq c(e_0^*, t)$ and $q_t(\tau=1, e_1) > q_t(\tau=0, e_0^*)$. Since $w > w_0 = u(0, t) - u(1, t)$, it follows that the agent gets strictly higher value under configuration $(\tau=1, e_1)$ than under configuration $(\tau=0, e_0^*)$. This concludes the proof. ■

Proof of Fact 4: Indeed, $m_{G^{\pi,p}}(V) = k$ if and only if for all $k' \neq k$,

$$V\pi_k - p_k > V\pi_{k'} - p_{k'}. \quad (10)$$

For $k' < k$, this last condition is equivalent to $V \geq \max_{k' < k} \{(p_k - p_{k'}) / (\pi_k - \pi_{k'})\}$, which in turn is equivalent to $V > V_{k-1}$. Similarly, for $k' > k$, equation (10) is equivalent to $V_k > V$.

This concludes the proof. ■

References

- Angrist, Joshua D. and Guido W. Imbens**, “Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity,” *Journal of the American Statistical Association*, 1995, *90* (430).
- , – , and **Donald B. Rubin**, “Identification of Causal Effects using Instrumental Variables,” *Journal of the American Statistical Association*, 1996, *91* (434).
- Angrist, Joshua, Eric Bettinger, Erik Bloom, Elizabeth King, and Michael Kremer**, “Vouchers for Private Schooling in Colombia: Evidence from a Randomized Natural Experiment,” *American Economic Review*, December 2002, *92* (5), 1535–1558.
- Ashraf, Nava, Dean Karlan, and Wesley Yin**, “Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines,” *Quarterly Journal of Economics*, May 2006, *121* (2), 635–672.
- , **James Berry**, and **Jesse M. Shapiro**, “Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia,” *American Economic Review*, forthcoming.
- Banerjee, Abhijit**, “A Simple Model of Herd Behavior,” *The Quarterly Journal of Economics*, August 1992, *107* (3), 797–817.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak**, “Measuring Utility by a Single-Response Sequential Method,” *Behavioral Science*, 1964, *9* (3), 226–232.
- Berry, James, Greg Fischer, and Raymond Guiteras**, “Incentive Compatibility in the Wild: Field Implementation of the Becker-de Groot-Marshak Mechanism,” 2010. London School of Economics, *mimeo*.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch**, “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 1992, *100* (5), 992–1026.

- Bohm, Peter, Johan Lindén, and Joakin Sonnegård**, “Eliciting Reservation Prices: Becker-DeGroot-Marschak Mechanisms vs. Markets,” *The Economic Journal*, July 1997, 107 (443), 1079–1089.
- Chan, Tat Y. and Barton H. Hamilton**, “Learning, Private Information, and the Economic Evaluation of Randomized Experiments,” *Journal of Political Economy*, 2006, 114 (6), 997–1040.
- Cohen, Jessica and Pascaline Dupas**, “Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment,” *Quarterly Journal of Economics*, 2010, 125 (1), 1–45.
- Deaton, Angus**, “Instruments, Randomization, and Learning about Development,” *Journal of Economic Literature*, 2010, 48 (2), 424–455.
- Duflo, Esther, Michael Kremer, and Jonathan Robinson**, “How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya,” *American Economic Review*, 2008, 98 (2), 482–488.
- , **Rachel Glennerster, and Michael Kremer**, “Using Randomization in Development Economics Research: A Tool Kit,” in T. Paul Schultz and John Strauss, eds., *Handbook of Development Economics*, Vol. 4, Amsterdam: Elsevier, 2008, pp. 3895–3962.
- , **Rema Hanna, and Stephen Ryan**, “Monitoring Works: Getting Teachers to Come to School,” 2010.
- Dupas, Pascaline**, “Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya,” 2009. University of California, Los Angeles *mimeo*.
- , “What Matters (and What Does Not) in a Households’ Decision to Invest in Malaria Prevention,” *American Economic Review*, 2009, *forthcoming*.
- , “Short-Run Subsidies and Long-Term Adoption of New Health Products: Experimental Evidence from Kenya,” 2010. University of California, Los Angeles *mimeo*.
- **and Jonathan Robinson**, “Savings Constraints and Microenterprise Development: Evidence from a Field Experiment in Kenya,” 2009. University of California, Los Angeles, *mimeo*.
- Flood, A.B., J.E. Wennberg, R.F. Nease, F.J. Fowler, J. Ding, and L.M. Hynes**, “The Importance of Patient Preference in the Decision to Screen for Prostate Cancer,” *Journal of General Internal Medicine*, 1996, 11 (6), 342–349.
- Gertler, Paul**, “Do Conditional Cash Transfers Improve Child Health? Evidence from PROGRESA’s Control Randomized Experiment,” *American Economic Review*, 2004, 94 (2), 336–341.

- Heckman, James J. and Edward Vytlacil**, “Structural Equations, Treatment Effects, and Econometric Policy Evaluation,” *Econometrica*, May 2005, *73* (3), 669–738.
- Imbens, Guido W.**, “Better LATE Than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009),” 2010. Harvard University, *mimeo*.
- **and Joshua D. Angrist**, “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, March 1994, *62* (2), 467–475.
- Jadad, Alejandro R. and Murray Enkin**, *Randomized Controlled Trials: Questions, Answers, and Musings*, BMJ Books, 2007.
- Jin, Hui and Donald B. Rubin**, “Principal Stratification for Causal Inference with Extended Partial Compliance,” *Journal of the American Statistical Association*, 2008, *103* (481), 101–111.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” 2009. University of Chicago, *mimeo*.
- Karlan, Dean S. and Jonathan Zinman**, “Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment,” *Econometrica*, 2010, *forthcoming*.
- Keller, L. Robin, Uzi Segal, and Tan Wang**, “The Becker-DeGroot-Marschak Mechanism and Generalized Utility Theories: Theoretical Predictions and Empirical Observations,” *Theory and Decision*, 1993, *34* (2), 83–97.
- King, Michael, Irwin Nazareth, Fiona Lampe, Peter Bower, Martin Chandler, Maria Morou, Bonnie Sibbald, and Rosalind Lai**, “Impact of Participant and Physician Intervention Preferences on Randomized Trials: A Systematic Review,” *Journal of the American Medical Association*, 2005, *293* (9), 1089–1099.
- Kremer, Michael and Edward Miguel**, “The Illusion of Sustainability,” *The Quarterly Journal of Economics*, 2007, *122* (3), 1007–1065.
- , – , **and Rebecca Thornton**, “Incentives to Learn,” *The Review of Economics and Statistics*, 2009, *91* (3), 437–456.
- Miguel, Edward and Michael Kremer**, “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities,” *Econometrica*, January 2004, *72* (1), 159–217.
- Milgrom, Paul and Ilya Segal**, “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 2002, *70* (2), 583–601.
- **and John Roberts**, “Price and Advertising Signals of Product Quality,” *The Journal of Political Economy*, 1986, *94* (4), 796–821.

- **and** – , “Relying on the Information of Interested Parties,” *The Rand Journal of Economics*, 1986, 17 (1), 18–32.
- Moffitt, Robert**, “Estimating Marginal Treatment Effects in Heterogeneous Populations,” 2008. Johns Hopkins University, *mimeo*.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer**, “Coarse Thinking and Persuasion,” *Quarterly Journal of Economics*, 2008, 123 (2), 577–619.
- Nguyen, Trang**, “Information, Role Models and Perceived Returns to Education: Experimental Information, Role Models and Perceived Returns to Education: Experimental Evidence from Madagascar,” 2009. MIT, *mimeo*.
- Oster, Sharon M.**, *Strategic Management for Nonprofit Organizations: Theory and Cases*, Oxford, UK: Oxford University Press, 1995.
- Philipson, Tomas and Jeffrey Desimone**, “Experiments and Subject Sampling,” *Biometrika*, 1997, 84 (3), 619–631.
- **and Larry V. Hedges**, “Subject Evaluation in Social Experiments,” *Econometrica*, 1998, 66 (2), 381–408.
- Plott, Charles R. and K. Zeiler**, “The Willingness to Pay–Willingness to Accept Gap, The ‘Endowment Effect,’ Subject Misconceptions, and Experimental Procedures for Eliciting Valuations,” *American Economic Review*, 2005, 95 (3), 530–545.
- Rayo, Luis and Ilya Segal**, “Optimal Information Disclosure,” 2008. Stanford University, *mimeo*.
- Rothschild, Michael**, “A Two-Armed Bandit Theory of Market Pricing,” *Journal of Economic Theory*, 1974, 9 (2), 185–202.
- Roy, A.D.**, “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 1951, 3 (2), 135–146.
- Scharfstein, Daniel O., Andrea Rotnitzky, and James M. Robins**, “Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models,” *Journal of the American Statistical Association*, 1999, 94 (448), 1096–1120.
- Schultz, T. Paul**, “School Subsidies for the Poor: Evaluating the Mexican Progresa Poverty Program,” *Journal of Development Economics*, 2004, 74 (1), 199–250.
- Silverman, W.A. and D.G. Altman**, “Patients’ Preferences and Randomised Trials,” *The Lancet*, 1996, 347 (8995), 171–174.
- Stolberg, Harald O., Geoffrey Norman, and Isabelle Trop**, “Randomized Controlled Trials,” *American Journal of Roentgenology*, 2004, 183 (6), 1539–1544.

- Thornton, Rebecca**, “The Demand for and Impact of Learning HIV Status: Evidence from a Field Experiment,” *American Economic Review*, 2008, *98* (5), 1829–1863.
- Tilbrook, Helen**, “Patients’ Preferences within Randomised Trials: Systematic Review and Patient Level Meta-analysis,” *BMJ*, 2008, *337*, 1864–1871.
- Volpp, Kevin G., Andrea Gurmankin Levy, David A. Asch, Jesse A. Berlin, John J. Murphy, Angela Gomez, Harold Sox, Jingsan Zhu, and Caryn Lerman**, “A Randomized Controlled Trial of Financial Incentives for Smoking Cessation,” *Cancer Epidemiology Biomarkers & Prevention*, 2006, *15* (1), 12.
- Volpp, K.G., L.K. John, A.B. Troxel, L. Norton, J. Fassbender, and G. Loewenstein**, “Financial Incentive-based Approaches for Weight Loss: A Randomized Trial,” *Journal of the American Medical Association*, 2008, *300* (22), 2631–2637.
- Zelen, Marvin**, “A New Design for Randomized Clinical Trials,” *New England Journal of Medicine*, 1979, *300* (22), 1242–1245.