Online Appendix for: Selective Trials: A Principal-Agent Approach to Randomized Controlled Experiments

By Sylvain Chassang, Gerard Padró i Miquel, and Erik Snowberg

Proofs

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, with infinite samples, $G_0 \preceq G$.*

PROOF:

The data $\mathbf{d}_G$ can be broken 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 $i \in \mathbb{N}$, set $m_i^h = \emptyset$, $p_i^h = 0$, and draw τ_i^h as a Bernoulli variable of parameter π_0. Finally, set $y_i^h = y_{\sigma_{\tau_i^h}(i)}$. It is easy to check that indeed, $h(\mathbf{d}_G) \sim \mathbf{d}_{G_0}$.

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

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

PROOF:

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

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

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 $h : \mathcal{D} \to \Delta(\mathcal{D})$ 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 $\mathbb{N} \to \mathbb{N}$. 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 $h : \mathcal{D} \to \Delta(\mathcal{D})$. It is easy to check that $h(\mathbf{d}_G) \sim \mathbf{d}_{G'}$.[1] This concludes the proof.

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.*

PROOF:

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.

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

PROOF:

Agents of type t and t' are such that $V_t > V_{t'}$ and $m_G(t) \neq m_G(t')$. Denote $\pi(m) = \text{Prob}(\tau = 1|m)$ and $p_{m_G} = \mathbb{E}_{\mu_G(\cdot|m)}[p]$. By optimality of the message, it must be that

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

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

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

[1]Note 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.

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

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

PROOF:

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

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

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 (optimally) sent by a type with value V_1, and m_2 (optimally) sent by a type with value V_2. Let $p_{G'}(m) = \mathbb{E}_{\mu_{G'}(\cdot|m)}[p]$. We must have that

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

within mechanism G'. These two inequalities yield that $(\pi_{G'}(m_2) - \pi_{G'}(m_1))V_1 \le p_{G'}(m_2) - p_{G'}(m_1) \le (\pi_{G'}(m_2) - \pi_{G'}(m_1))V_2$, which implies that $|p_{G'}(m_2) - p_{G'}(m_1)| < (\overline{\rho} - \underline{\rho})V_{\max}$. 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 $|(\pi_{G'}(m_1) - \pi_{G'}(m_2))V - p_{G'}(m_1) + p_{G'}(m_2)| \le 2(\overline{\rho} - \underline{\rho})V_{\max}$.

PROPOSITION 4 (most informative mechanisms): *Any strictly incentive-compatible blind mechanism G identifies at most the 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'.

PROOF:

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

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 $\varphi \in (0, 1)$, $V_t(\phi{=}\varphi) < \varphi V_t(\phi{=}1)$.

PROOF:

The proof is given for the general case where there might be multiple optimal effort choices. Let $V_t(\tau, e)$ denote the expected value of type t under treatment status τ and when expending 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 $\arg\max_{e\in E} V_t(\tau=1,e) \cap \arg\max V_t(\tau=0,e) \neq \emptyset$, the inequality is an equality and, since we normalized $V_t(\phi=0)=0$ we obtain that $V_t(\varphi)=\varphi V_t(\phi=1)$. Inversely, if $\arg\max_{e\in E} V_t(\tau=1,e) \cap \arg\max V_t(\tau=0,e) = \emptyset$, the inequality is strict and $V_t(\varphi) < \varphi V_t(\phi=1)$.

PROPOSITION 6 (identifying perceived returns to effort): *For any value φ,*

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

PROOF:
The result follows directly from applying the Envelope Theorem to (1).

PROPOSITION 7 (identifying perceived success rates):

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

PROOF:
The result follows directly from applying the Envelope Theorem to (2).

FACT 3: *Assume that outcome $y=1$ yields strictly greater utility than $y=0$, that is, $u(y=1,t) > u(y=0,t)$, and an agent perceives treatment to be beneficial:*

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

$$\textit{Then,} \quad w_{0,t} = \max\{w \mid V_t(\tau=1,w) = V_t(\tau=0,w)\}.$$

PROOF:
Whenever $w = w_{0,t}$, 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 $w > w_{0,t}$, $V_t(\tau=1,w) > V_t(\tau=0,w)$. 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 $w >$

$w_{0,t} = u(0,t) - u(1,t)$, 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.

FACT 4: *Under mechanism $G^{\pi,p}$ an agent of type t sends message k if and only if $V_t \in [V_{k-1}, V_k]$.*

PROOF:

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

$$V\pi_k - p_k > V\pi_{k'} - p_{k'}. \tag{A1}$$

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 $k' > k$, (A1) is equivalent to $V_k > V$. This concludes the proof.

A Numerical Example

This section illustrates the step-by-step 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.

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 expenditure. 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 down the set of possible values.[2]

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

[2]For simplicity, we consider priors that put point masses on a few possible states. Unfortunately, such strong priors often 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 our somewhat unusual parameter values.

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

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.[3] 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 expend much effort into using the water treatment product more thoroughly $(R_e = 1/2, R_{e,t} = 1/4.)$.

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 $(R_b,\ R_e,\ R_{e,t})$. Based on this, we illustrate the six possible values of V_t in the following matrix:

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

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, as 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.

[3]For 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 (Rothschild, 1974; Banerjee, 1992; Bikhchandani et al., 1992).

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 $\frac{\partial V_t(\tau,w)}{\partial w}\Big|_{w=0} = R_{b,t} + R_{e,t}^2$ and $\frac{\partial \widehat{\Delta}(w)}{\partial w}\Big|_{w=0} = R_e \times R_{e,t}$. With this data the principal can identify:

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

and thus, the rest of the unknown parameters: $R_e = \frac{\partial \widehat{\Delta}(w)}{\partial w}\Big|_{w=0} \Big/ R_{e,t}$, $R_{b,t} = \frac{\partial V_t(\tau,w)}{\partial w}\Big|_{w=0} - R_{e,t}^2$, $R_b = \widehat{\Delta} - R_e \times R_{e,t}$. The same information can be identified in a mathematically simpler, but more data intensive, way by identifying $w_{0,t}$ 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 a principal 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 agent's indirect preferences over the water treatment product, and bonuses associated with staying healthy, allows the principal to infer the agents' beliefs about the effects of casual and more thorough usage. This, in turn, allows the principal 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.

*

REFERENCES

Banerjee, Abhijit, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, August 1992, *107* (3), 797–817.

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.

Rothschild, Michael, "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 1974, *9* (2), 185–202.