A New Perspective on the Issue of Selection Bias into Randomized Controlled Field Experiments

Michèle Belot and Jonathan James

No. 23 /14

BATH ECONOMICS RESEARCH PAPERS

Department of Economics

A New Perspective on the Issue of Selection Bias into Randomized Controlled Field Experiments

Michèle Belot
University of Edinburgh

Jonathan James
University of Bath

May 2, 2014

Abstract

Most randomized controlled trials require the collaboration of field partners. Such self-selection could introduce a potential bias, because only the most optimistic may participate. We revisit this prediction. We argue that in many situations, the experimental intervention is competing with alternative interventions participants could conduct themselves outside the experiment. Since participants have a chance of being assigned to the control group, participating has a direct opportunity cost, which is likely to be higher for optimists. We propose a model of self-selection and show that both pessimists and optimists may opt out of the experiment, leading to an ambiguous selection bias.

Keywords: Field experiments, selection bias, randomised controlled trials, external validity

1 Introduction

The last decade has seen a booming increase in popularity of field experiments in economics and social sciences (Holt, 2005) and there is now a stronghold of researchers advocating the case for randomized controlled trials (RCTs) to overcome the challenges of endogeneity plaguing empirical work (Burtless, 1995, Duflo and Kremer, 2005). One key issue though with such experiments is the issue of external validity, specifically, we would like to know if we can generalize the lessons we draw from such experiments to the wider population of interest.[1] The answer to this question depends on the representativeness of the sample along characteristics that may be correlated with the estimated treatment effects. In that context, self-selection deserves prime attention. In almost all randomized

[1] A number of recent papers study self-selection in field experiments: Allcott and Mullainathan (2012), Belot and James (2013), Gautier and van der Klaauw (2012)

controlled experiments in social sciences, participation is voluntary. Often RCTs require the collaboration of field practitionners such as job agencies, schools, firms or charities who are prepared to collaborate with researchers and provide the necessary support for data collection. On top of that, the population to be treated (job seekers, pupils, employees or donors for example) often also has the right to opt out from participating to the experiment.[2]

Researchers are well aware of the implications of self-selection for external validity. The standard prediction in the literature on selection in randomised field experiments is a positive selection bias, following the spirit of the literature on policy evaluation (Heckman and Vytlacil, 2006 and Heckman et al.,1999). List and Rasul (2011) state in a recent review paper on "Field Experiments in Labor Economics" that "*Indeed, in almost any social experiment related to job training programs, it is a concern that those most likely to benefit from the program select into the program.*"

We argue that this prediction is not necessarily correct because participating to the experiment does not guarantee treatment and entails a chance of ending up in the control group. Being assigned to the control group often entails an *opportunity cost*, in the sense that participants are often required not to conduct any competing intervention (and certainly not the experimental intervention itself) at the same time. Almost all RCTs (in medical and social sciences) entail such opportunity cost. Our own experiment (Belot et al. (2013)) provides an example. We study the process of selection of schools in a randomized controlled experiment aimed at testing the effects of two incentive schemes (a piece rate and a competition) to encourage children to eat more fruit and vegetables. We assigned schools to either of the two incentives or to a control group, which would not implement any incentive scheme over the course of the experiment. We approached schools and explained the randomization procedure. It turned out that some schools explicitly asked whether they could be in the treatment group (which of course we did not allow) and a number of schools who opted out explicitly mentioned to us that they preferred not to participate because they were already involved in similar programmes. Their decision not to participate could therefore introduce a *negative* selection bias as well. Similarly, a recent study by Fryer (2011) tests the effects of incentives in schools

[2]Many Ethical Review Boards require explicit informed consent from participants to experiments. This means that participants must be told in advance what treatment is tested and how the randomisation will take place.

and randomize schools into various types of incentive schemes. The schools in the control group did not implement any incentive scheme over the course of the study.

This paper re-visits the issue of self-selection in randomized field experiments. We present a simple model of self-selection allowing for the presence of competing interventions. Potential partners self-select based on a prior about the effectiveness of the intervention. We introduce two specific features in the model, likely to be relevant for field experiments in social sciences: First, participating to a randomized controlled trial entails a probability of not receiving the treatment. Second, the alternative to participating is to conduct an intervention competing with the intervention to be tested with the RCT. Thus, being in the control group entails an opportunity cost and is not necessarily equivalent to not participating to the experiment. Those who choose not to participate could conduct interventions that are (1) similar to the intervention proposed (e.g. an incentive scheme) and/or (2) directly competing with the intervention proposed, in the sense that both cannot be implemented at the same time (such as again in the example of a specific incentive scheme).

In many situations, the intervention to be tested already exists and the main goal of the RCT is to establish causality. The key reason why causality is difficult to establish is because of self-selection: Those who are already exposed to the treatment are not a random sample of the population and are likely to be those who benefit most from the treatment. In fact, researchers often look for "virgin samples", i.e. participants who have not yet been exposed to the treatment. Or, alternatively, participants themselves will not see the point of participating if they are already implementing the intervention. For example, a firm that implements a tournament incentive scheme may not be interested in testing the effects of a tournament, particularly if it means there is a chance they may have to give it up temporarily (or have a smaller proportion of their workers exposed to the tournament scheme) if the end up in the control group. Of course this obviously introduces a selection bias which, suprisingly, has received little attention in the literature.

In order to illustrate this point, we provide in Appendix A a brief overview of the information provided in field experimental studies published over the last five years in the top 5 journals and in the American Economic Journal: Applied Economics. We focus on the fields of policy evaluation, personnel economics and development economics, which have

all experienced a significant increase in the popularity of field experimental research.[3] In most cases the experimental interventions cannot be considered as completely innovative, i.e. they are not interventions that participants could not have considered themselves before, and an opportunity cost is likely to be present. Typically the samples used had not been treated before, although in some cases this was not alluded to or it was not clear. Thus, the typical experiment tests an intervention that is plausibly accessible outside the experiment. Or at least similar interventions are accessible to the potential participants. As we will see, this has important implications for the prediction of the direction of the bias.

Our work relates most to Malani's work (2008) on self-selection in medical randomised controlled trials. His model resembles the one presented here, to the extent that participants self-select into the experiment based on a prior about the effectiveness of the treatment. He assumes that patients have a choice between an old treatment and participating to a randomised controlled trial (RCT) involving a new treatment. As the probability of being assigned to the treatment group in the RCT increases, less optimistic patients are willing to participate to the RCT. If the probability is lower than 1 (which will always be the case by definition in a RCT), only the most optimistic will be willing to experiment and we will obtain a positive selection bias. The main difference with our model is that we allow for a possible positive correlation in effectiveness between the experimental intervention and the alternative intervention that participants could implement outside the experiment. In fact, it could even be that the experimental intervention itself is directly available to the participants outside the experiment. This contrasts sharply with medical trials, where the new treatment to be tested in not available to the participants outside the RCT.

The question of selection bias also has a specific angle in a policy relevant context. Policymakers may not be interested in the average treatment effect calculated over the entire population. Instead, they may be interested in knowing whether the intervention is worth subsidizing, whether it is sufficiently effective for those who are "on the margin" and would not conduct it without subsidy. We find in that case, it is possible to sign the selection bias if one carefully documents the costs of experimentation beforehand.

The remaining of the paper is structured as follows. Section 2 presents the model

[3]We searched these journals systematically for the keywords 'field' and 'experiment'.

of selection. Section 4 discusses the nature of the bias for the policy-relevant treatment effect. Section 4 concludes.

2 The Model

The goal of an experiment is to evaluate the effectiveness of a particular intervention on an outcome of interest and within a population of interest. The objective of this section is to present a simple model that sheds light on the nature of selection in such environments.

2.1 The traditional problem of selection bias

Researchers are interested in estimating the causal effect of a treatment on an outcome of interest. Suppose we are interested in the average treatment effect, i.e. that is the treatment effect averaged over the population of interest. For example, in Belot et al. (2013), we are interested in testing whether providing incentives can encourage children to develop healthier diet habits. Let us define y as the outcome of the intervention (for example, the number of fruit and vegetables consumed in a week).

Suppose the researcher is interested in testing the effectiveness of a specific treatment/intervention, say T_A, with $T_A = 0$ if the individual is not treated and $T_A = 1$ if the individual is treated. The level y achieved by individual i is a function of the treatment:

$$y_i(T_{Ai}) = \beta_{0i} + \beta_{Ai} T_{Ai} + u_i, \tag{1}$$

where β_{0i} and β_{Ai} are unknown parameters and u_i is an individual specific error term. Obviously there is an issue of selection bias only if there is heterogeneity in the effectiveness of an intervention across individuals (β_{Ai} are individual specific) in a way that cannot be fully controlled for (i.e. there are conditional on unobservable characteristics of the population). In general terms β_{Ai} is a function of observable and unobservable characteristics X_i : $\beta_{Ai} = g(X_i, u_i)$.

Let us ignore for now the possibility of conducting an RCT. If the intervention was available to everyone, the decision to conduct the intervention would depend on the expected marginal benefit of the intervention. More precisely, suppose individuals have a prior β^*_{Ai} regarding the effectiveness of treatment A. β^*_{Ai} is a draw from a distribution $F(\beta_{Ai})$,and assume that $E(\beta^*_{Ai}) = \beta_{Ai}$ (without loss of generality)[4] and defined on the

[4]Note that the argument carries through even if priors are systematically biased upwards or down-

support $[\beta_{A,\min}, \beta_{A,\max}]$. Suppose conducting the intervention has a fixed cost c_A that we assume identical across individuals.

The expected net benefit of conducting the intervention is $\beta^*_{Ai} - c_A$. Only those with $\beta^*_{Ai} > c_A$ will conduct the intervention. This type of selection corresponds to the "traditional" selection bias usually considered in the policy evaluation literature. Those who would implement the treatment are those with the highest expected marginal benefits of the treatment. Given that priors are correlated with the truth, the estimated average treatment effect $\overline{\widehat{\beta_A}}$ will be a positively biased estimate of $\overline{\beta_A}$.

2.2 Randomized-Controlled Experiments

2.2.1 Unbiased estimate

The main point of an RCT is to get rid of the selection problem described above. Suppose the researcher is interested in estimating $\overline{\beta}_A$ (the average treatment effect across all individuals). An unbiased estimate of $\overline{\beta}_A$would be obtained by a randomised controlled trial, with a group of N_{Treat} individuals assigned at random to a treatment group (who receive treatment A) and a group of $N_{Control}$ individuals assigned at random to a control group (who do not receive the treatment). The key assumption to obtained an unbiased estimate of $\overline{\beta}_A$ is $E(T_{Ai}|u_i) = 0$, that is the assignment to treatment and control is random and is uncorrelated with unobservable characteristics. In that case:

$$\begin{aligned} \overline{\widehat{\beta}}_A &= \frac{\sum_{i \in Treatment} y_i}{N_{Treat}} - \frac{\sum_{i \in Control} y_i}{N_{Control}} \\ &= \overline{y}_{Treat} - \overline{y}_{Control} \end{aligned} \tag{2}$$

2.2.2 Outside Option and Competing interventions

We now discuss the issue of self-selection in a world where possible *competing* and *similar* interventions may exist. For simplicity and without loss of generality, let us assume there is one possible alternative intervention T_B.

Definition 1. *Intervention T_B is said to be competing with the experimental intervention T_A if it cannot be (practically) implemented in conjunction with the experimental intervention.*

For example, the firm that implements a tournament cannot implement a piece-rate for the same workers.

wards.

Definition 2. *Intervention T_B is said to be similar to the intervention T_A if the prior beliefs of the treatment effects of interventions T_A and T_B $(\beta^*_{A_i}, \beta^*_{B_i})$ are positively correlated.*

For example, schools could implement various interventions to increase fruit and vegetable consumption, and could hold positively correlated beliefs regarding the effectiveness of these interventions.

A special case is if T_A itself is available outside the experimental setting.

An important question is what does participating in the experiment imply for the participants. Specifically, what can the control group do? What interventions can they implement? And how does this compare to what non-participants can do?

In principle, those in the control group are *not treated* and may be involved in other interventions. There are two typical cases. First, participants in the control group are told not to implement T_A or any other treatment that is both competing and similar to T_A. Second, which is typical in medical trials, the control group can continue to take the current best treatment that is available (T_B), but will not have access to T_A. If that is the case, then the estimated effect will only capture the net difference between the effect of the experimental intervention and the effect of alternative competing interventions. In contrast, non-participants can always implement T_B or even T_A if it is available to them.

2.2.3 Participation Decision and Selection Bias

In this section we will derive the participation constraints for the two cases described above. We consider first the situation where $E(\beta^*_A) \geq E(\beta^*_B)$ and $E(\beta^*_A) \geq 0$, that is, the experimental intervention is a priori believed to be more effective than the interventions that are currently available, and T_A is believed to have a positive effect.

Presumably there is a direct cost of participating to the experiment, which we denote k (such as providing support for data collection, etc.) but also a potential subsidy s. Also, implementing intervention T_B has a cost c_B (assumed identifical for all individuals).

Denote $\bar{g} = \pi(s_{treat} - k - c_A) + (1 - \pi)(s_{control} - k)$.

Case 1 Intervention T_B is not available to the control group.
Then the decision to participate to the experiment must satisfy:

$$\pi\beta^*_{Ai} + \bar{g} \geq \max\{\beta^*_{Bi} - c_B, 0\}, \tag{3}$$

where β^*_{Bi} is i's prior belief about the effectiveness of T_B.

We now derive the participation constraints. We first start with the optimists (such that $\beta^*_{Bi} \geq c_B$). The decision to participate to the experiment satisfies:

$$\begin{aligned} \pi\beta^*_{Ai} + \bar{g} &\geq \beta^*_{Bi} - c_B \qquad \text{if} \quad \beta^*_{Bi} \geq c_B \quad (4) \\ &\Leftrightarrow \\ \beta^*_{Bi} &\leq \pi\beta^*_{Ai} + \bar{g} + c_B \qquad \text{if} \quad \beta^*_{Bi} \geq c_B \quad (5) \end{aligned}$$

Without loss of generality, let us take the example where β^*_{Bi} is a linear projection of β^*_{Ai}, i.e. $\beta^*_{Bi} = \rho\beta^*_{Ai} + \eta_i$, with ρ being a fixed parameter and η_i an error term, with mean zero and variance σ^2_η. We can then write condition [4] as:

$$(\rho - \pi)\beta^*_{Ai} \leq \bar{g} + c_B - \eta_i \qquad \text{if} \quad \beta^*_{Bi} \geq c_B \quad (6)$$

If $\rho > \pi$:

$$\beta^*_{Ai} \leq \frac{\bar{g} + c_B - \eta_i}{(\rho - \pi)} \qquad \text{if} \quad \beta^*_{Bi} \geq c_B \quad (7)$$

If $\rho > \pi$, then there will be a negative selection bias. All else equal, the higher the correlation between T_A and T_B the more negative selection there is.

If $\rho < \pi$, then positive selection could take place:

$$\beta^*_{Ai} \geq \frac{\bar{g} + c_B - \eta_i}{(\rho - \pi)} \qquad \text{if} \quad \beta^*_{Bi} \geq c_B \quad (8)$$

In contrast, the participation constraint for pessimists (those such that $\beta^*_{Bi} < c_B$) satisfies:

$$\begin{aligned} \pi\beta^*_{Ai} + \bar{g} &\geq 0 \qquad \text{if} \quad \beta^*_{Bi} < c_B \quad (9) \\ &\Leftrightarrow \quad (10) \\ \beta^*_{Ai} &\geq -\frac{\bar{g}}{\pi} \quad (11) \end{aligned}$$

These two conditions show that positive and negative selection can take place at the same time. Without further assumptions on the distribution and joint distribution of β^*_{Ai} and β^*_{Bi}, we cannot draw conlusions on the direction of the bias.

Case 2 Intervention T_B is available to the control group.

This second case corresponds to Malani's model and implies a lower implicit cost of participating to the experiment. But more importantly, the implications for selection are very different.

Here the decision to participate to the experiment must satisfy:

$$\begin{aligned}
\pi\beta^*_{Ai} + (1-\pi)\max\{\beta^*_{Bi} - c_B, 0\} + \bar{g} &\geq \max\{\beta^*_{Bi} - c_B, 0\}, && (12)\\
\pi\beta^*_{Ai} + (1-\pi)(\beta^*_{Bi} - c_B) + \bar{g} &\geq \beta^*_{Bi} - c_B \quad \text{if} \quad \beta^*_{Bi} - c_B \geq 0 && (13)\\
\pi(\beta^*_{Ai} - \beta^*_{Bi} + c_B) + \bar{g} &\geq 0 \quad \text{if} \quad \beta^*_{Bi} - c_B \geq 0 && (14)\\
\pi(\beta^*_{Ai} - \beta^*_{Bi}) &\geq -c_B - \bar{g} \quad \text{if} \quad \beta^*_{Bi} - c_B \geq 0 && (15)
\end{aligned}$$

If $\beta^*_{Ai} > \beta^*_{Bi}$, which we have assumed is true in expectations, then those who are optimistic about A will opt in, then the constraint is always satisfied if $c_B + \bar{g} > 0$, which is likely to be the case in a typical randomised controlled experiment. In that case, no negative selection takes place.

However, if the non participants have access to T_A, while the control group does not (and only has access to T_B), then the participation condition becomes:

$$\begin{aligned}
\pi\beta^*_{Ai} + (1-\pi)(\beta^*_{Bi} - c_B) + \bar{g} &\geq \beta^*_{Ai} - c_A \quad \text{if} \quad \beta^*_{Ai} - c_A \geq \beta^*_{Bi} - c_B \geq 0\\
\beta^*_{Ai} &\leq \frac{c_A + \bar{g}}{(1-\pi)} + (\beta^*_{Bi} - c_B)\\
(1-\rho)\beta^*_{Ai} &\leq \frac{c_A + \bar{g}}{(1-\pi)} - c_B + \eta_i,
\end{aligned}$$

such that the most optimistic will select out to implement T_A outside the experiment.

There could still be positive selection in the case where B is not attractive to non participants or participants in the control group.

$$\begin{aligned}
\pi\beta^*_{Ai} + \bar{g} &\geq 0 \quad \text{if} \quad \beta^*_{Bi} - c_B < 0,\\
\beta^*_{Ai} &\geq -\frac{\bar{g}}{\pi} \quad \text{if} \quad \beta^*_{Bi} - c_B < 0.
\end{aligned}$$

3 Selection bias and policy-relevant average treatment effect estimate

While researchers may be interested in estimating the value of $\overline{\beta_A}$ defined over the entire population, policy makers may not be interested in that parameter. Instead, they

may be more interested in what we will call a "policy-relevant average treatment effect" (PRATE)[5]. We now define the PRATE and discuss the implications of selection.

Suppose a policy maker considers implementing an intervention at a given point in time. Suppose the social optimum is such that it is beneficial to implement the intervention if $\beta_{Ai} \geq c_A - e$, where e is a positive externality on society, which we assume is known. Think for example of the reduced costs of health care associated with a healthy diet.

Suppose the policy maker would like to identify the average effect on the population for which the intervention is socially optimal but not privately optimal. In that case, the experiment would ideally participate to the experiment excludes (1) those with $\beta_{Ai} \geq c_A$, who would implement the intervention without public intervention, (2) those for whom $\beta_{Ai} < c_A - e$. We will refer to the population corresponding to this interval $[c_A - e, c_A]$ as the "target" population.

Since the *distribution* of treatment effects is unknown, the goal of a policy experiment is to obtain an estimate of the average treatment effect in the target population. Of course an alternative could be to simply elicit beliefs from the population. But this alternative is unlikely to be implementable with compatible incentives and, moreover, beliefs are noisy and the goal of an RCT could be to obtain more precise estimates of the treatment effects.

Suppose the intervention is assessed according to its average costs and benefits and is judged worthwhile if the average benefits are higher than the average costs in the target population. Then the Policy Relevant Average Treatment Effect corresponds to the treatment effect estimated for the target population.

If that is the case and if beliefs are not systematically biased, then the experimental sample that would provide an unbiased estimate of this average policy relevant treatment effect is defined on that same interval.

Definition 3. *The policy relevant average treatment effect* β^P *is defined as* $\beta^P = \int_{c_I - e}^{c_I} \beta_{Ai} dF(\beta_{Ai})$

Thus at the top of the distribution, the experimental sample should include those such that $\beta^*_{Ai} \leq c_A$ and those such that $\beta^*_{Ai} \geq c_A - e$.

The question is whether it is possible to choose a subsidy level that would lead to the appropriate self-selection, that is, a level of subsidy such that only individuals in the

[5]Note that our concept of PRATE is different from the policy relevant treatment effect discussed in Heckman, J.J. and E. Vytlacil (2001)

target population would self-select.

For simplicity, we focus on the most interesting and simple case where T_A is available outside the experiment, T_A is more attractive than T_B and T_B is not available to the control group. Policy makers (or experimental researchers) can choose freely a level of subsidy that will compensate the participants both in the control and treatment groups. Let us denote $\overline{s} = \pi s_{treat} + (1-\pi)s_{control}$

There are two parameters the experiment can control: π and $\overline{s}$.

$$\begin{aligned} \pi\beta^*_{Ai} + \overline{s} - k - \pi c_A &\geq \max\{\beta^*_{Ai} - c_A, 0\} && (16) \\ &\Leftrightarrow \\ \beta^*_{Ai} &\leq c_A + \frac{\overline{s}-k}{(1-\pi)} \qquad \text{if} \quad \beta^*_{Ai} \geq c_A && (17) \\ \beta^*_{Ai} &\geq c_A - \frac{\overline{s}-k}{\pi} \qquad \text{if} \quad \beta^*_{Ai} < c_A && (18) \end{aligned}$$

To ensure that the relevant experimental sample participates (defined on the interval $[c_A - e, c_A]$), the following conditions must be satisfied

$$\begin{aligned} \frac{\overline{s}-k}{(1-\pi)} &= 0 \\ \frac{\overline{s}-k}{\pi} &= e \end{aligned}$$

It is obvious that there exists no combintion of π and $\overline{s}$ that could achieve this support. In fact, as soon as the expected subsidy $\overline{s} > k$, then both types of selection take place at the same time. As soon as the average subsidy covers more than the experimental costs k, the experimental sample will for sure include participants for which the intervention is privately desirable (i.e. such that $\beta^*_{Ai} \geq c_A$).

But if $\frac{\overline{s}-k}{\pi} \leq e$, then we now for sure that the experimental sample will not include participants for whom the intervention is neither privately nor socially desirable. That is, we now for sure that the PRATE will be biased up upwards.

4 Conclusion

This paper discusses the implications of self-selection into randomised controlled field experiments. We point out that in many situations, alternatives to the experimental intervention (or the intervention itself) are available outside the experiment. The implication is that being part of the control group entails an opportunity cost, which could lead to both positive and negative selection at the same time.

We also discuss the potential bias in what we call the "policy-relevant treatment effect", which corresponds to the treatment effect for those individuals for whom the treatment is socially efficient. We derive conditions under which the bias can be signed and, more specifically, is unambiguously positive.

References

Allcott, H. and S. Mullainathan (2012), External Validity and Partner Selection Bias, *NBER Working Papers* 18373

Bandiera, O., I. Barankay and I. Rasul (2007), Incentives for Managers and Inequality Among Workers: Evidence From a Firm-Level Experiment, Quarterly Journal of Economics 122(2), 729-773.

Belot, M. and J. James (2013), Partner Selection into Policy Relevant Field Experiments, *mimeo*

Belot, M., J. James and P. Nolen (2013), Incentives and Children's Dietary Choices: A Field Experiment in Primary Schools, *mimeo*

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

Crépon, B., E. Duflo, M. Gurgand, R. Rathelot and Philippe Zamora (2013), Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment, *The Quarterly Journal of Economics* 128 (2): 531-580

Duflo, Esther & Michael Kremer (2005). "Use of Randomization in the Evaluation of Development Effectiveness." in George Pitman, Osvaldo Feinstein and Gregory Ingram

(eds.) Evaluating Development Effectiveness. New Brunswick, NJ: Transaction Publishers.

Duflo, E., Dupas, P., & M. Kremer (2011), Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya, American Economic Review 101(5), 1739-1774

Dupas, P. (2011), Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya, American Economic Journal: Applied Economics 3(1), 1-34.

Fehr, E. & L. Goette (2007), Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment, American Economic Review 97(1), 298-317.

Fryer, R. Jr. (2011), Financial Incentives and Student Achievement: Evidence from Randomized Trials, *The Quarterly Journal of Economics* 126 (4): 1755-1798

Gautier, P. and B. van der Klaauw (2012), Selection in a field experiment with voluntary participation, *Journal of Applied Econometrics* 27, 63-84.

Hastings, J. and J. Weinstein, School Choice, and Academic Achievement: Evidence from Two Experiments (2008), *Quarterly Journal of Economics*, 123(4), 1373-1414

Heckman, J.J. and E. Vytlacil (2001), Policy-Relevant Treatment Effects, *The American Economic Review* 91(2), 107-111

Heckman, James and E.J. Vytlacil (2006), Econometric Evaluation of Social Programs. *Handbook of Econometrics*, Volume 6, James Heckman and E. Leamer, eds., Amsterdam: North Holland.

Heckman, James, R. Lalonde and J. Smith (1999), The Economics and Econometrics of Active Labor Market Programs, *Handbook of Labor Economics*, Volume 3, Ashenfelter, A. and D. Card, eds., Amsterdam: Elsevier Science.

Holt, Charles, A. (2005), *Markets, Games and Strategic Behavior: Recipes for Interactive Learning.* Addison-Wesley, Boston

List, John A. & Rasul, I. (2011), Field Experiments in Labor Economics, *Handbook of Labor Economics*, Elsevier .

Malani, A. (2008), Patient enrollment in medical trials: Selection bias in a randomized experiment, *Journal of Econometrics* 144(2), 341-351, June.

Wisdom, J., J.S. Downs and G. Loewenstein (2010), Promoting Healthy Choices: Information versus Convenience, American Economic Journal: Applied Economics, 2(2), 164-78.

Appendix A: Survey of Top 5 field experimental papers

Paper	Documentation of partner selection	Comparison with broader sample	Similar & competing interventions available	Virgin sample	Innovative intervention
Fryer, QJE (2011) Partner: 200 schools across three cities in the US Intervention: Provision of incentives to school children (on school inputs and outputs)	Yes	Yes	Yes	Yes	Yes
Hastings and Weinstein, QJE (2008) Partner: Charlotte Mecklenburg Public School Distict Intervention: Provision of information about school outcomes to parents	No	No	Yes	Yes (cohort) No (school)	No
Wisdom, Downs, and Loewenstein, AEJ:Applied (2010) Partner: Fast Food Sandwich Chain Intervention: Provision of nutritional information	No	No	Yes	Yes	No
Dupas, AEJ:Applied (2011) Partner: two rural districts of Western Kenya, involving 328 primary schools Intervention: Provision of HIV risk information to teenagers	No	No	Yes	Unclear	No
Duflo, Dupas and Kremer, AER (2011) Partner: Schools in Kenya Intervention: Tracking in schools	No	No	Yes	Yes	No
Duflo, Dupas and Kremer, AER (2011) Partner: Four schools in Chitwan District, Nepal Intervention: Providing sanitary products on school attendance	No	No	Yes	Unclear	No
Bandiera, Barankay and Rasul, QJE (2007) Partner: fruit picking farm Goal: Test the effectiveness of worker's pay incentive scheme	No	No	Yes	Yes	No
Fehr and and Goette, AER (2007) Partner: Bicycle Messenger Service Goal: Testing labour supply responses to transitory wage changes	No	No	Yes	Yes	No