

The Encouragement Design for Program Evaluation

17 September 2007

Alexis Diamond (IFC & Harvard University)
Jens Hainmueller (Harvard University & IFC)

Summary

This short essay describes a methodology for program evaluation known as encouragement design that is particularly useful when one wishes to rigorously estimate the effect of some intervention (such as accessing financing or simplified business registration) that cannot itself be randomly administered to some and not others. The method is called an encouragement design because it requires that a randomly-selected group of beneficiaries receive extra encouragement to undertake the intervention. Encouragement typically takes the form of additional information or incentives. A successful encouragement design allows estimation of the effect of the intervention as well as the effect of the encouragement itself. IFC is currently planning an encouragement design in an Alternate Dispute Resolution (Business Enabling Environment) project in Macedonia.

Experimental Design

A randomized experiment is a study designed to estimate the effect of an intervention on a set of indicators ("results") for a collection of units (e.g., individuals or firms). In a randomized experiment, the researcher can manipulate the exposed group (the "treatment group") and the not-exposed group (the "control group"). Randomization is useful because it ensures that (on average), prior to the intervention, treatment and control groups are essentially identical and therefore would show very similar results in the absence of the treatment. Therefore, a difference in results for the two groups can be causally attributed to the program.

The Encouragement Design

In a classic experiment, the researcher has full control over which subjects will be exposed to an intervention. In practice, experiments can be difficult to implement because it is difficult to ensure that all those selected for treatment get the treatment and all selected for control do not. Also, it is sometimes impractical or unethical to conduct experiments and force some of the subjects to be in the control group. For example, in an access-to-finance program, we would like to know the effect of receiving a loan on firm-level impacts, but we cannot randomize loan approvals. Similarly, in a business-enabling-environment program we would like to estimate the effect of registering a business on firm-level impacts, but we cannot force some firms to become registered and other firms to remain informal.

The *encouragement design* is a special case of an experimental design that can be used in situations with little control over subjects' compliance. The key idea is that instead of randomizing the application of the intervention itself, what is randomized is *encouragement* to receive the treatment. By randomizing encouragement and carefully tracking outcomes for all those who do and do not receive the encouragement, it is possible to obtain reliable estimates of BOTH the encouragement and the intervention itself (see "Quantities of Interest", below).

Encouragement may take the form of information that is additional to whatever is already part of program implementation and targeted at the unit (e.g., firm) level. For example, in the context of a business-registration simplification program, the project team may plan to advertise the simplification on TV and radio; the

encouragement could take the form of additional direct mailings to a random sample of firms. In the case of access-to-finance, encouragement can take the form of loan-application training—the advantage here is that one can then estimate the effect of training on receiving a loan as well as the effect of the loan on firm-level impacts.

Note that encouragement is merely that—encouragement. Some firms receiving encouragement may not follow through with their loan or business registration application. And some firms who do not receive encouragement may get loans or register their businesses. All that is required is that the encouragement increases the likelihood that units will follow through with what they are being encouraged to do.

Quantities of Interest

A. The effect of encouragement

If encouragement (e.g., information, training, etc.) is randomly applied to some units and not others, then estimating its effect on any measurable result is straightforward: simply compare results for the randomly-selected encouraged group vs. results for the randomly-selected not-encouraged group. This quantity of interest, known as the “Intention-to-Treat” effect, or ITT, is the effect of the encouragement itself, e.g., “What is the effect of training on applying for a loan?” “What is the effect of direct-mail advertising on business registration?”

Since encouragement itself is randomized, the comparison between the encouraged (e.g., trained) and not-encouraged groups will be free of any bias due to self-selection if encouragement was delivered as planned. What is important here is that one can measure results for all subjects from the sample in which randomized selection was undertaken. Whenever encouragement takes the form of key project activities (e.g., training, or firm-level informational campaigns), measuring the impact of encouragement is particularly useful.

B. The effect of the treatment

In many situations we are also interested in the effect of the larger event or intervention (the “treatment”, e.g., obtaining the financing, registering the business) on firm-level impacts. We wish to answer questions such as: “What is the effect of accessing finance on employment and revenues?” This effect is often difficult to estimate because up-take of treatment itself is not randomized and therefore we cannot simply compare results for treated vs. untreated firms.

However, one can still estimate the effect of the treatment by exploiting a randomized encouragement design, adjusting the ITT effect by the amount of non-compliance. This yields to the local average treatment effect (LATE), computed as:

$$\text{LATE} = \text{ITT} / \text{Compliance Rate}$$

where Compliance Rate = Fraction of Subjects that were treated in the treatment group - Fraction of Subjects that were treated in the control group

If the compliance rate is 100%, LATE = ITT, we have perfect compliance, and all assigned to the treatment take the treatment and all those assigned to the control do not take the treatment. The compliance rate can be thought of as the fraction of subjects that fall into the sub-population of “compliers”, the group for whom the decision to take treatment was directly affected by the assignment. Put differently, this is the group induced by the encouragement to take advantage of the treatment.

Notice that the compliers can be thought of as the group of people that actually stick to the experimental protocol; they will take the treatment if assigned to the treatment group, but they will not take the treatment if assigned to the control group. From the point-of-view of a policymaker, compliers are an interesting sub-population because they are the only ones who are actually affected by the encouragement. Note that not all subjects in the sample will be compliers: some will always follow through with treatment, and some will never take the treatment regardless of their assignment.

Usually, the compliance rate will be less than one, and then it becomes important to recognize that the LATE effect estimates the effect of treatment only for the sub-population of compliers and it does not constitute the effect of the treatment for the whole sample.

An important special case is when the control group can be excluded from taking the treatment. Then non-compliance can only occur in the treatment group and the LATE will be equal to the average treatment effect for the treated—that is, the average effect of the treatment for those that do take the treatment. In general, the compliance rate depends on the encouragement. Some encouragements will be relatively effective and strongly influence take up. Other encouragements may be less effective.

There are a few important assumptions that need to hold for the LATE to give an unbiased estimate.

- *Encouragement cannot have perverse consequences; it cannot make subjects less likely to receive the treatment.* This is often a reasonable assumption, but needs to be considered carefully on a case-by-case basis.
- *Encouragement is in fact randomly assigned; those assigned to encouragement receive it, and the rest do not.*
- *Encouragement has no direct effect on results, except via increasing the probability of receiving treatment.* This is why encouragement should be kept as simple as possible. For example, if encouragement takes the form of a training session that could have its own effect on results *in addition to* encouraging take-up of treatment, then this assumption would be violated. Such a case would be a training session that combines lessons on accessing finance *and* successful business management.