Can Participants Estimate Program Impacts? Evidence from a Labor Market Intervention

Andrew Brudevold-Newman, Maddalena Honorati, Pamela Jakiela, and Owen Ozier*

August 17, 2017

Abstract

We test whether participants in two labor market interventions hold accurate beliefs about program impacts, exploiting a randomized evaluation to compare participants' beliefs to intent-to-treat estimates. We provide a framework for comparing methods of belief elicitation and identifying potential biases. We find the program participants do a poor job of estimating their own counterfactual (probabilistic) outcomes. However, we extend the existing set of best practices by demonstrating that participants are quite good at estimating average treatment impacts on the population once behavioral biases are taken into account.

Keywords: occupational choice, youth unemployment, beliefs, impact evaluation

*Brudevold-Newman: University of Maryland, abn@umd.edu; Honorati: Social Protection and Labor Global Practice, World Bank, mhonorati@worldbank.org; Jakiela: University of Maryland and IZA, pjakiela@umd.edu; Ozier: Development Research Group, World Bank, oozier@worldbank.org. We are grateful to Sarah Baird, Maya Eden, David Evans, Deon Filmer, Jessica Goldberg, Markus Goldstein, Joan Hamory Hicks, David Lam, Isaac Mbiti, David McKenzie, Patrick Premand, seminar participants at Duke, USC, and the University of Oklahoma, as well as numerous conference attendees for helpful comments. Rohit Chhabra, Emily Cook-Lundgren, Gerald Ipapa, and Laura Kincaide provided excellent research assistance. This research was funded by the IZA/DFID Growth and Labour Markets in Low Income Countries Programme, the CEPR/DFID Private Enterprise Development in Low-Income Countries Research Initiative, the ILO's Youth Employment Network, the National Science Foundation (award number 1357332), and the World Bank (SRP, RSB, i2i, Gender Innovation Lab). We are indebted to staff at the International Rescue Committee (the implementing organization) and Innovations for Poverty Action for their help and support. The findings, interpretations and conclusions expressed in this paper are entirely those of the authors, and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments of the countries they represent.

1 Introduction

Impact evaluation is an increasingly critical part of development aid. Many aid agencies and donors now explicitly favor "evidence-based" policies that have been shown to have significant impacts on participants, and randomized trials and other impact evaluations account for an increasing share of academic research on international development. Given the tremendous lengths one must go to in order to produce credible estimates of a program's impacts, an important question is whether participants themselves understand the effects of the programs in which they participate. It is not uncommon for labor market programs to survey participants *ex post*; however, Smith, Whalley, and Wilcox (2012) find that such *ex post* assessments of a program's impact are not highly correlated with objective measures of program effects. Understanding participants' beliefs about program impacts is important for two reasons. Most obviously, if — through their participation — participants obtain reasonable estimates of program impacts, this information may be a feasible, low-cost alternative to formal impact evaluation. On the other hand, if program participants do not understand a program's impacts, even after they have participated in the program, it is hard to imagine that they are making optimal decisions about whether or not to participate.

We test whether participants in two labor market interventions hold accurate beliefs about program impacts, exploiting a randomized evaluation to compare participants' beliefs to intent-to-treat estimates. Brudevold-Newman, Honorati, Jakiela, and Ozier (2017) evaluate a a "microfranchising" program that offered young women in some of Nairobi's poorest neighborhoods a combination of vocational and life skills training together with start-up capital and ongoing business mentoring, comparing applicants randomly assigned to the franchise treatment to both a control group and a comparison group that received an unrestricted cash grant. They find that both the franchise treatment and the cash grant had significant and durable impacts on the likelihood of self-employment, although impacts on income disappeared in the second year after treatment.

In this paper, we study the beliefs of participants randomly assigned to the franchise and grant treatments. Building on recent work by Smith, Whalley, and Wilcox (2012) and McKenzie (2016), we show that participants do a poor job of estimating their own counterfactual probabilities of self-employment and paid work. However, an alternative approach to belief estimation that asks participants to estimate frequencies of different outcomes in reference populations of (hypothetical) individuals similar to them generates extremely accurate estimates of average treatment effects.

2 Empirical Approach

As Smith, Whalley, and Wilcox (2012) point out, one reason participant evaluations of programs may differ from rigorous estimates of program impacts is that participant evaluation questions are often quite open-ended. For example, participants in the National Job Training Partnership Act program were asked "Do you think that the training or other assistance that you got from the program helped you get a job or perform better on the job?" (Smith, Whalley, and Wilcox 2011, p. 9). This question is obviously problematic because it is not at all clear whether better on-the-job performance should be linked to any measurable outcome (e.g. income); moreover, the link between the fraction of participants who believe that the program had a positive impact and the estimated treatment effect of the program is unclear, making it difficult to test whether participants' subjective evaluations are accurate. Smith, Whalley, and Wilcox (2012) suggest replacing such subjective evaluation questions with alternatives that (i) clearly specify the outcomes and time periods of interest, (ii) ask for continuous (as opposed to binary) responses that can be directly compared to ITT estimates, and (iii) make the counterfactual nature of the question transparent.

We follow the recommendations of Smith, Whalley, and Wilcox (2012) and ask participants in the franchise and grant treatments to estimate the counterfactual probabilities of self-employment and paid work for a reference group of women similar to themselves. Specifically, we ask women in each of the two treatment arms the question: "I would like you to imagine 100 women from [your neighborhood] who applied to the [name of treatment arm] program but who were not admitted into it. In other words, please think about 100 women similar to yourself who were not selected to the [name of treatment arm] program. Out of 100 women, how many do you think are currently running or operating their own business?" We also ask an analogous question about involvement in paid work for others. Smith, Whalley, and Wilcox (2012) suggest using this question to construct a perceived counterfactual, which can then be compared with the average outcome in the treatment group. We take a different approach, asking each participant to estimate how many of 100 women similar to themselves who "applied for and were admitted into" the program were (at the time of the survey) operating their own business (and, in a subsequent question, we ask how many were doing paid work for others). We calculate each participant's belief about the treatment effect of the program (on, for example, self-employment) by taking the difference between the perceived frequency of self-employment among women invited to participate in the program and the perceived frequency of self-employment among similar women who were not invited to participate.

We also test a second method proposed by Smith, Whalley, and Wilcox (2012): asking participants about the probability that they would be self-employed (or doing paid work for others) in the absence of the program. These individual-level beliefs about one's own counterfactual can then be combined with data on actual outcomes to construct estimates of perceived treatment effects. However, as Smith, Whalley, and Wilcox (2012) emphasize, there are several drawbacks to this approach. First, program participants may find it inherently difficult to imagine what their lives would have been like in the absence of the program. For example, psychological studies of "hindsight bias" suggest that people have a difficult time remembering the beliefs they held in the past and tend to assume that realized outcomes were always foreseeable (Fischhoff 1975, Madarász 2012). In our context, we might expect that those who have received vocational training and gained self-employment experience might have a difficult time remembering that they had not always known how to operate a business; thus, hindsight bias might inflate participants' estimates of their own counterfactual, particularly among successful microentrepreneurs. Estimates of one's own counterfactual may also be biased by the tendency to attribute one's own success to individual agency as opposed to external factors (Miller and Ross 1975). This would lead those who have benefited from business or vocational training to overstate the likelihood that they would have started a successful business in the absence of the program.

In the context of our evaluation, a third problem with questions designed to elicit beliefs about one's own counterfactual probability of self-employment (or paid work) is that they are unlikely to work well when respondents have low levels of numeracy. Though almost 92 percent of the women in our sample completed primary school, a relatively large number are not familiar with the concept of percentages. Roughly one in four cannot (correctly) answer the question: "If there is a 75 percent chance of rain and a 25 percent chance of sun, which type of weather is more likely?" While it is possible to elicit probabilistic expectations from subjects with no prior knowledge of probability, it is costly and time-consuming to do so. Instead, we asked every subject categorical questions about their counterfactual probabilities of self-employment and paid work, and collected more specific data on counterfactual probabilities from those who successfully answered the screening question described above.[1]

[1] We worded the categorical question to make responses directly comparable to probability estimates. Respondents chose one of the following options: (1) *In the absence of the program, I would definitely be self-employed*, (2) *In the absence of the program, I would probably be self-employed but it is not certain*,

3 Framework for Interpreting Empirics

To facilitate comparisons between different approaches to belief elicitation, we introduce a simple conceptual framework that formalizes the measurement issues highlighted above. First, consider an outcome, y, and a program whose causal effect on that outcome is to increase its expected value by $\beta > 0$. Let γ denote the expected value of y in the absence of the program: $E[y_j|T_j = 0] = \gamma$.

We wish to know whether program participants hold accurate beliefs about β. Let

$$\tilde{\beta}_i = \tilde{\beta} + \phi_i \tag{1}$$

denote participant i's belief about the impact of the program, and let

$$\tilde{E}[y_j|T_j = 0] = \tilde{\gamma} + \nu_i \tag{2}$$

be participant i's belief about the expected value of the outcome of interest for an untreated individual j who is outwardly similar to her. $\tilde{\beta}$ is the average belief about the impact of the program, and $\tilde{\gamma}$ is the average belief about the outcome of interest in the eligible population in the absence of the program. ϕ_i is the idiosyncratic component of beliefs about the impact of the program; without loss of generality, we assume that the distribution of ϕ_i is mean zero, and we let σ_ϕ denote its variance. ν_i can be decomposed into a mean-zero error term and a term which reflects the perceived difference between the population average of y and one's own counterfactual:

$$\nu_i = \tilde{\alpha_i} \cdot \mathbb{1}(j = i) + \epsilon_i. \tag{3}$$

As discussed above, asking participants about their own counterfactuals may be problematic

(3) *In the absence of the program, the chances of me being self-employed or not self-employed are equal*, (4) *In the absence of the program, I would probably not be self-employed but it is not certain*, or (5) *In the absence of the program, I would definitely not be self-employed.*

(for example, because of hindsight bias), and the population mean of these $\tilde{\alpha}_i$ values, $\tilde{\alpha} = E[\tilde{\alpha}_i]$ may not be equal to 0.[2] Combining and generalizing these expressions, respondents report:

$$\tilde{E}[y_j|T_j] = \tilde{\beta} \cdot T_j + \tilde{\gamma} + \tilde{\alpha}_i \cdot \mathbb{1}(j = i) + \phi_i \cdot T_j + \epsilon_i \tag{4}$$

Specifically, when asked to report the rate of self-employment among 100 potential program participants who were not invited to participate in the program, a respondent in our study reports:

$$\tilde{E}[y_j|T_j = 0] = \tilde{\gamma} + \epsilon_i. \tag{5}$$

When asked to report the rate of self-employment among 100 potential program participants who were invited to participate in the program, she reports:

$$\tilde{E}[y_j|T_j = 1] = \tilde{\beta} + \tilde{\gamma} + \phi_i + \epsilon_i. \tag{6}$$

Finally, when asked to report her own counterfactual probability of self-employment, a participant reports:

$$\tilde{E}[y_i|T_i = 0] = \tilde{\gamma} + \tilde{\alpha}_i + \epsilon_i. \tag{7}$$

The framework presented above helps to clarify the distinctions between the different approaches to estimating participant beliefs. First, consider an estimate of participant beliefs constructed by taking the average belief about one's own counterfactual (in our context, the counterfactual probability of self-employment) and subtracting this from the observed outcome in the treatment group. The expected value of this estimator is:

$$\begin{aligned} E[y_j|T_j = 1] - E[\tilde{E}[y_i|T_i = 0]] &= \beta + \gamma - (\tilde{\gamma} + \tilde{\alpha} + E[\epsilon_i]) \\ &= \beta + (\gamma - \tilde{\gamma}) - \tilde{\alpha} \end{aligned} \tag{8}$$

[2] This may be thought of as a "Lake Wobegon" effect.

since $E[\epsilon_i] = 0$. Thus, this estimator will be biased if participants hold inaccurate beliefs about the counterfactual probability of self-employment, and it will be biased when psychological factors such as hindsight bias lead participants to overstate their own counterfactual probability of self-employment. The second estimator proposed by Smith, Whalley, and Wilcox (2012) is constructed by subtracting the mean rate of self-employment in a reference group of untreated women from the observed rate of self-employment in the treatment group. The expected value of this estimator is given by:

$$\begin{aligned} E[y_j|T_j = 1] - E[\tilde{E}[y_j|T_j = 0]] &= \beta + \gamma - (\tilde{\gamma} + E[\epsilon_i]) \\ &= \beta + (\gamma - \tilde{\gamma}) \end{aligned} \tag{9}$$

This estimator overcomes the behavioral issues inherent in estimating one's own counterfactual. However, when estimates of participant beliefs constructed in this manner diverge from actual program impacts, it is impossible to determine whether participants hold inaccurate beliefs about the impact of the program or inaccurate beliefs about the counterfactual.

The outcomes of interest in impact evaluations are often difficult to measure, and considerable effort goes into the design and pre-testing of questionnaires. Nonetheless, there is no guarantee that outcome measures derived from survey questions (for example, about labor market participation) and participant responses to belief-elicitation questions will line up, particularly in low-income settings where formal, full-time employment is relatively uncommon (and there is continuous variation in the number of hours worked, and labor supply varies substantially from week to week).[3] Impact evaluation questions designed to

[3] Smith, Whalley, and Wilcox (2012) are aware of this issue and recommend asking extremely specific questions: for example, what fraction of participants meet a well-specified criterion for employment — for example, working more than 35 hours per week — which can then be used to construct the empirical estimate of the programs impact. However, such precisely worded questions are not always feasible. In our context, we worried that any question of the form "Out of 100 women, how many spend at least X hours operating their own business?" would be substantially more difficult to answer than a less specific question because few people work full-time and there is no obvious break in the distribution of hours worked at any

measure beliefs about the counterfactual may reveal systematic deviations between participants' beliefs about outcome levels and actual outcome levels; however, such measurement error is only problematic if it cannot be separated from the quantity of interest. To address this issue, we propose an estimate of participant beliefs that is calculated by taking the difference between beliefs about the mean outcome of interest in a reference population of treatment versus control individuals:

$$\begin{aligned} E[\tilde{E}[y_j|T_j = 0]] - E[\tilde{E}[y_j|T_j = 0]] & \\ &= \tilde{\beta} + \tilde{\gamma} + E[\phi_i] + E[\epsilon_i] - (\tilde{\gamma} + E[\epsilon_i]) \\ &= \tilde{\beta} \end{aligned} \tag{10}$$

Such an estimator allows for a direct test of the hypothesis that participants hold accurate beliefs about program impacts; moreover, collection of the relevant data necessarily also allows researchers to assess the related issue of whether participants can estimate the counterfactual — allowing for a comparison of the different approaches of belief estimation.

4 Results

Our results, which are summarized in Figure 1, suggest that participants hold remarkably accurate beliefs about program impacts. The figure compares ITT estimates of program impacts to estimates of participant beliefs about program impacts calculated by taking the difference in reference group probabilities for the treatment and control groups.[4] For example, the ITT estimates suggest that the franchise treatment increased the likelihood of self-employment by 11.9 percentage points; those assigned to the program believe

point.

[4] In other words, beliefs were estimated by asking women assigned to each treatment group to estimate reference group probabilities (frequencies) for both the treatment and comparison groups. Women assigned to the control group were not asked to estimate a reference group probability for those assigned to the treatment groups since they were not familiar with the details of each treatment.

that it increased the likelihood of self-employment by 12.3 percentage points. Similarly, those assigned to the cash grant treatment believe that it increased the likelihood of self-employment by 10.6 percentage points; the ITT estimates suggest a 12.9 percentage point increase. Those assigned to the franchise treatment also have remarkably accurate beliefs about the program's impact on the likelihood of paid employment. Those assigned to the cash grant treatment have less accurate beliefs about the program's impact on paid employment, though they are appropriately signed and well within the confidence interval of the estimated treatment effect. Thus, our results suggest that participants' do a reasonably good job of estimating the impact of programs that they have participated in. For the outcome most directly impacted by the treatments (self-employment), participants do a remarkably good job of estimating the program's impacts.

Figure 2 compares beliefs about the probability of self-employment and paid work to levels observed in the treatment and control groups, and compares beliefs about one's own counterfactual to beliefs about a reference population of untreated women. Several patterns are apparent. First, women in the franchise treatment group underestimate the probability of paid work in both the treatment and the control group. Consequently, an estimate of the impact of the franchise program on the probability of paid work that compared counterfactual beliefs to observed levels in the treatment group would perform very poorly. Women in both the franchise and grant treatments hold more accurate beliefs about the level of self-employment (in both the treatment and control groups); however, women in both treatment arms seem to overestimate the frequency of self-employment and underestimate the frequency of paid work in both the treatment and the control groups. Thus, differences between observed outcome levels and participant beliefs appear to be systematic, suggesting that it will typically be better to estimate program beliefs by comparing beliefs about the control group to beliefs about the treatment group (rather than the observed outcome levels in the treatment group).

The figure also demonstrates that concerns that estimates of one's own counterfactual might be biased appear well-founded: the average of own counterfactual estimates is consistently higher than the estimated outcome for a reference population of untreated women. This pattern is particularly pronounced for the franchise treatment, most dramatically when participants are asked to report their own counterfactual probability of self-employment. Though participants hold accurate beliefs about the level of self-employment in both the treatment and control groups, own counterfactual estimates are so inflated that they suggest a negative impact of the program on self-employment. Thus, our evidence clearly supports the view that own counterfactual estimates are of little use in estimating treatment effects. This finding is consistent with recent work by McKenzie (2016); he finds that program participants (business owners) do a very poor job of estimating the counterfactual. Our results support his conclusion, but suggest that an alternative approach to eliciting participants' beliefs performs substantially better.

5 Conclusion

Individual beliefs are a critical component in any model of occupational choice, and are therefore an important object of study in their own right. If job-seekers cannot correctly estimate the impacts of training programs and other active labor market interventions, they cannot make optimal decisions about human capital investments and occupational choice. We study beliefs directly and show that program participants do a very poor job of estimating their own counterfactual probabilities of self-employment and paid work. However, we provide a methodological improvement: we asked participants to estimate the counterfactual for people other than themselves. With this refinement, they do correctly estimate the overall impact of two labor market interventions on occupational choice. Our evidence also illustrates a mechanism by which people can hold seemingly inconsistent beliefs about policies: they can accurately believe the policy helps someone else, while inaccurately believing it did not help them. In the context of elections, for example, this could lead people to vote against their own interests.

References

Brudevold-Newman, A., M. Honorati, P. Jakiela, and O. Ozier (2017): "A Firm of One's Own: Experimental Evidence on Credit Constraints and Occupational Choice," World Bank Policy Resarch Working Paper No. 7977.

Fischhoff, B. (1975): "Hindsight Is Not Equal to Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty.," *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 288.

Madarász, K. (2012): "Information Projection: Model and Applications," *The Review of Economic Studies*, 79(3), 961–985.

McKenzie, D. (2016): "Can Business Owners Form Accurate Counterfactuals? Eliciting Treatment and Control Beliefs about Their Outcomes in the Alternative Treatment Status," World Bank Policy Research Working Paper 7768.

Miller, D. T., and M. Ross (1975): "Self-Serving Biases in the Attribution of Causality: Fact or Fiction?," *Psychological Bulletin*, 82(2), 213.

Smith, J., A. Whalley, and N. Wilcox (2011): "Are Program Participants Good Evaluators?," working paper.

——— (2012): "Are Participants Good Evaluators?," working paper.

Figure 1: Participants' Beliefs about Impacts of Treatments

Panel A: Beliefs about Impact of Franchise Treatment

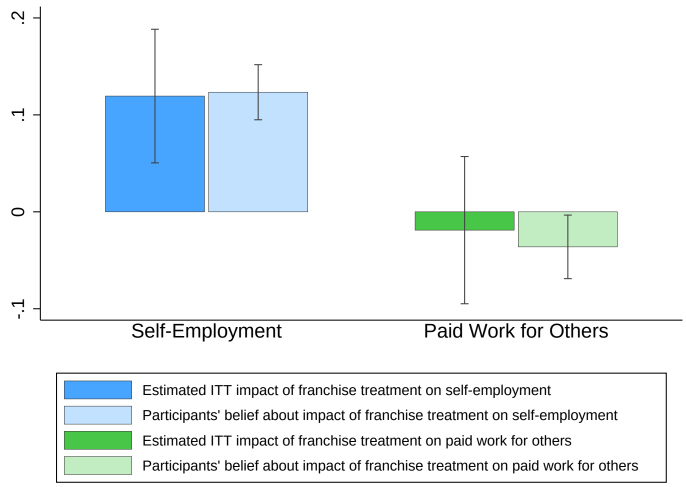

Panel B: Beliefs about Impact of Grant Treatment

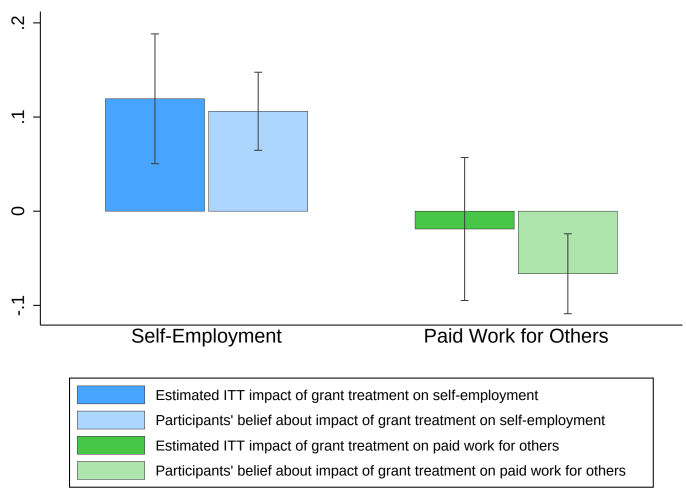

ITT estimates of treatment are estimated via OLS, controlling for stratum fixed effects (we omit other controls included in our main specifications to make ITT estimates as comparable to self-reported beliefs as possible, though these controls have minimal impacts on estimated coefficients). Beliefs are estimated using estimates of the frequency of outcomes in a reference class of young women similar to oneself. For example, the estimate of the impact of the franchise treatment on the probability of self-employment is constructed using average responses to two questions: (1) "I would like you to imagine 100 women from [your neighborhood] who applied to the [name of treatment arm] program and were admitted into it, just as you were. In other words, please think about 100 women similar to yourself. Out of 100 women, how many do you think are currently running or operating their own business?" and (2) "Now I would like you to imagine 100 women from [your neighborhood] who applied to the [name of treatment arm] program and but who were not admitted into it. In other words, please think about 100 women similar to yourself who were not selected to the [name of treatment arm] program. Out of 100 women, how many do you think are currently running or operating their own business?" The difference in responses to these two questions (divided by 100) is the individual-level estimate of the average treatment effect of the program on self-employment.

Figure 2: Participants' Beliefs about Impacts of Treatments

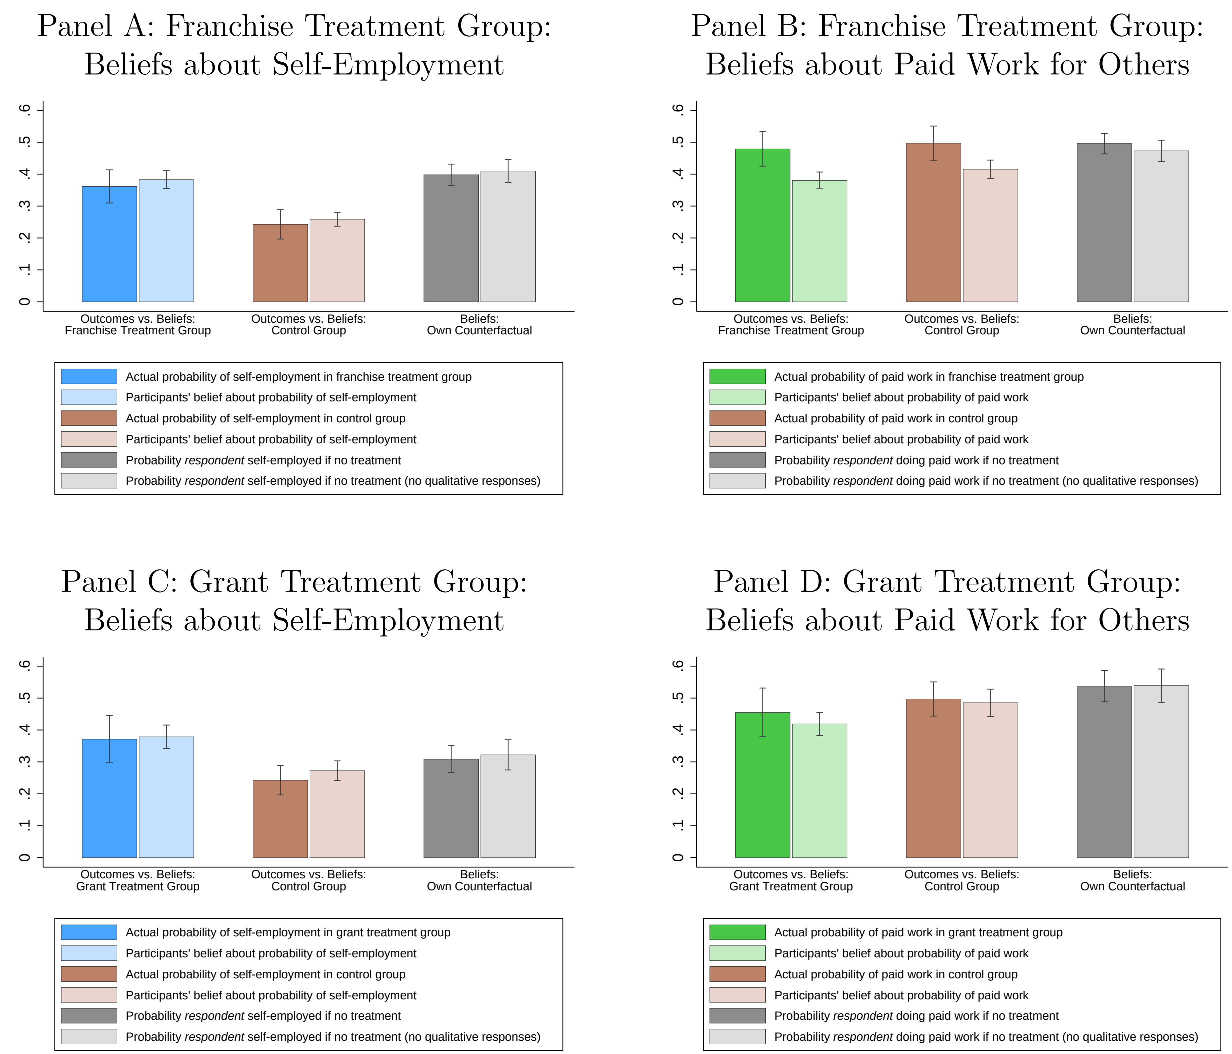

The figure compares observed levels of self-employment and paid work in the treatment groups and the control group to beliefs about levels held by women assigned to the franchise and grant treatment arms. See Figure 1 for a description of the belief elicitation questions. The probability that a respondent would be doing paid work or in self-employment in the absence of treatment is the average response to a question about the counterfactual likelihood of involvement in the labor market.