ABDUL LATIF JAMEEL POVERTY ACTION LAB

CASE STUDY 4: VOCATIONAL TRAINING IN COLOMBIA

Threats & Analysis


Man working in a bakery near Barranquilla, Colombia. Photo: Paul Smith J-PAL/IPA.

This case study is based on two papers: Attanasio, Orazio, Arlen Guarin, Carlos Medina, and Costas Meghir. 2017. "Vocational Training for Disadvantaged Youth in Colombia: A Long-Term Follow-up." *American Economic Journal: Applied Economics* 9(2): 131-143. Attanasio, Orazio, Adriana Kugler, and Costas Meghir. 2011. "Subsidizing Vocational Training for Disadvantaged Youth in Colombia: Evidence from a Randomized Trial." *American Economic Journal: Applied Economics* 3 (3): 188-220.

J-PAL thanks the authors for allowing us to use their papers as a teaching tool.

KEY VOCABULARY	
Treatment assignment	An individual's treatment assignment is the group to which they were randomly assigned (a treatment or comparison group).
Treatment status	An individual's treatment status is what actually happened to them: were they treated or not?
Selection bias	Selection bias occurs when individuals who receive or opt into the program are systematically different from those who do not. Consider an elective after-school tutoring program. Is it effective at raising children's exam scores? Comparing scores for those who participate and those who don't will produce a biased estimate of the effect of the tutoring program if these groups differ across characteristics that correlate with test scores. For example, those who choose to participate may be more motivated, and may have scored better than non-participants even without the tutoring program. Randomization minimizes selection bias because it breaks the link between characteristics of the individual and their treatment status. Selection bias can occur in other ways in a randomized evaluation. For example: ● Participants can choose to take up a treatment or refuse it ● Participants can choose to leave the study (i.e., attrit)
Attrition bias	Attrition bias is a type of selection bias that occurs when people attrit or leave the study. This can bias the estimate of the treatment impact in two ways: 1. It may be the case that people with certain characteristics (e.g., those with the highest levels of education) in both the treatment and comparison groups leave. This means your study population looks less like the general population. The treatment effect you estimate might not represent the true effect for the general population. 2. The reasons people choose to leave may be correlated with the treatment. Suppose that the students with the highest grades in the after-school tutoring treatment group improve their performance and switch into elite private schools, leaving your study sample. Then comparing treatment and comparison groups after the program ends would underestimate the impact of the program, because the higher performing students are 'missing' from the treatment group.
Compliance	When a unit's treatment assignment (assigned to treatment or comparison group) matches their treatment status (took up or did not take up the program, respectively), we say they have complied. When a unit does not follow their treatment assignment, we have **non-compliance** at an individual/unit level. This means that at the study sample level we have partial compliance. Any study sample can be split into three distinct groups:

	1. **Compliers:** This group of people follows their treatment assignment. If they are assigned to the treatment group, they will take up the program; if they are assigned to the comparison group they will not take up the program. 2. **Always-takers:** This group of people will always take up the program, regardless of their treatment assignment. 3. **Never-takers:** This group of people will never take up the program, regardless of their treatment assignment. Note that we have the underlying assumption that there are no defiers, a fourth group who do the opposite of, or defy, their treatment assignment. Those individuals will not take up a program because they are assigned to the treatment group or will take it up because they are assigned to the comparison group.
Spillovers	Spillovers occur when the treatment indirectly affects those who have not been treated. Spillovers can be positive or negative.
Intention-to-Treat (ITT)	The ITT is a method of estimating the effect of the program that compares the average outcomes of those assigned to the treatment group to the average outcomes of those assigned to the comparison group, regardless of whether individuals within those groups have actually received the treatment. The ITT measures the impact of delivering a program in the real world, where some people don't take up the program when offered it, and others take up the program even when they are not expressly encouraged to do so. ITT = (avg. outcomes of those assigned to treatment) - (avg. outcomes of those assigned to comparison)
Local Average Treatment Effect (LATE)	The LATE is a method of estimating the effect of the program on those who complied with their treatment assignment (compliers). The LATE divides the ITT by the difference in the proportion of the treatment group who took up the program and the proportion of the comparison group who took up the program. Intuitively, you should think of the LATE as a way of adjusting the ITT to reflect that not all of those assigned to treatment were treated while some who were assigned to the comparison group were treated. $\text{LATE} = \dfrac{\text{ITT}}{\text{(proportion of take-up in treatment) - (proportion of take-up in comparison)}}$

LEARNING OBJECTIVE

This case study explores common threats to the validity of randomized evaluations and how they affect the estimation of a program's impact.

SUBJECTS COVERED

Attrition, selection bias, compliance, spillovers, intention-to-treat (ITT), local average treatment effect (LATE).

INTRODUCTION

In response to historically high youth unemployment in 1999-2000, the Colombian government launched the Jóvenes en Acción (Youth in Action), or JeA, program. The JeA program was a six-month vocational training program for young people between the ages of 18 and 25 in the bottom fifth of the income distribution. The JeA program provided three months of in-classroom training in a particular vocation (e.g., sewing, plumbing, cooking), a three month internship with a relevant company, and lessons on how to write a resume and apply for jobs.

The program was delivered by private vocational training institutions (Entidades de Capacitación, or ECAPs) located in cities and towns across the country. To enroll, young job-seekers were instructed to apply to the ECAP organization in their area. Each ECAP screened applicants to identify those who were most likely to benefit from their programs. The government compensated each ECAP based on how many young people successfully completed the internship with a participating firm, creating an incentive for ECAPs to screen out eligible applicants who were unlikely to succeed in the program.

Figure 1: Structure of the JeA program

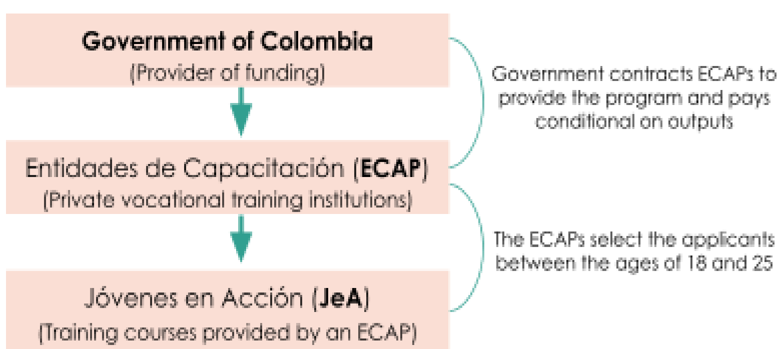

Researchers partnered with the ECAPs to conduct a randomized evaluation of the JeA program for the 2005 cohort. ECAPs identified eligible applicants through their screening process. Two-thirds of these screened applicants were randomly assigned to receive the vocational training program, and one-third were assigned to the comparison group. The research team conducted a baseline survey and follow-up telephone interviews 13 to 15 months after the conclusion of the program with a random sample of screened applicants to measure the program's short-term effects. They found large treatment effects on income and formal sector employment,[1] driven primarily by improvements among women.[2]

To measure the program's long-term impacts, researchers revisited the study sample in 2016. Due to the high costs of re-surveying the original sample, they instead matched the study sample to administrative data from Colombia's national database of contributions to health, pensions, and disability insurance, which captures anyone employed in the formal sector. They found the program had positive long-term effects on employment, earnings, and hours worked per week in the formal sector for both men and women.

This case study will take us through different threats to experimental validity, including spillovers, attrition, and partial compliance, and consider how they could influence the program's impact estimate. The discussion below largely draws from the JeA evaluation but also incorporates hypothetical examples that did not occur in the actual study.

THREATS TO VALIDITY IN RANDOMIZED EVALUATIONS

DISCUSSION TOPIC 1: SPILLOVERS

Spillovers occur when the outcomes of untreated units are indirectly affected by the treatment given to others. For example, when a parent vaccinates their child, that action also affects the health of neighbors' children, because they will now be slightly less likely to get sick.

In randomized evaluations, spillovers pose a challenge because they violate the key assumption that one unit's treatment assignment has no effect on the outcomes of others. In the case of immunizations, spillovers can make children in the comparison group *healthier* than they otherwise would be, leading us to underestimate the program's true effect.

1.1 In the case of the JeA program, can you think of any **positive spillovers**? Describe how they could happen.

[1] Formal economy is defined as the part of an economy that is regulated by government authorities through contract and company law, taxation and labor law. See UNESCWA (2022).
[2] There were some outcomes for which both men and women exhibited similar (and significant) treatment effects, namely the formality of work. Treated men and women were more likely to have formal employment, more likely to have a contract, and–conditional on being formally employed–have higher wages. The gender categories were limited to men and women at the time of the study.

1.2 Can you think of any potential **negative spillovers**? Describe how they could happen.

1.3 What are some strategies a research team could use to address potential spillovers? At what stage of the project should they be planned for and implemented?

DISCUSSION TOPIC 2: ATTRITION

Attrition occurs when study group members drop out of the study or data on them cannot be collected. This is a concern for several reasons. First, attrition—whether in the treatment or comparison group—reduces the sample size of the study, which makes it harder to detect the effect of the program. Second, attrition can cause bias. This bias can arise when certain types of people leave the study (e.g., those who live furthest from the village center or those from high-income households). If a specific type of person leaves the study in *both* the treatment and comparison group, then the study sample looks less like the general population, meaning the results of the study are harder to generalize. This affects the external validity of the study findings. More consequentially, if people with the certain characteristics leave the treatment or comparison group disproportionately, it reduces the balance of the two groups. This introduces bias into the estimate of the treatment effect and reduces the internal validity of the study findings.

To measure the JeA program's impact on employment outcomes, researchers conducted follow-up surveys in the short-term and linked participant records with administrative data from the Colombian government in the long-term. Attrition could be a concern in this context if we do not have updated contact information for participants or if we are not able to match participants with an administrative data record. In this section, we will examine the consequences of attrition for the evaluation of the JeA program.

Suppose there are 6,000 unemployed young people randomized into the treatment and comparison groups (4,000 in the treatment group and 2,000 in the comparison). Suppose all of the individuals assigned to the treatment group attended the JeA program, while none of the individuals assigned to the comparison group did. The formal employment status for members in each group under this scenario are shown for both baseline and endline in Table 1.

TABLE 1: EMPLOYMENT STATUS AT BASELINE AND ENDLINE				
	Baseline		Endline	
Employment status	Treatment	Comparison	Treatment	Comparison
Informal sector or unemployed	2,500	1,250	750	1,250
Formal sector	1,500	750	3,250	750
Sample Size	4,000	2,000	4,000	2,000

2.1 Using Table 1, calculate the following:

a. **At baseline**, what is the formal employment rate for each group?

b. **At endline**, what is the formal employment rate for each group?

c. What is the impact of the program on formal sector employment?

In the previous question we calculated the true effect of the intervention. Now we will see how different scenarios of attrition affect that estimate. Suppose now that in the comparison group, half of the jobseekers who remain without a formal job at the end of the year feel disillusioned and refuse to respond to the endline survey. The employment status of jobseekers in each group under this scenario is displayed in Table 2.

TABLE 2: EMPLOYMENT STATUS AT BASELINE AND ENDLINE WITH ATTRITION IN THE COMPARISON GROUP				
	Baseline		Endline	
Employment status	**Treatment**	**Comparison**	**Treatment**	**Comparison**
Informal sector or unemployed	2,500	1,250	750	625
Formal sector	1,500	750	3,250	750
Sample Size	4,000	2,000	4,000	1,375

2.2 Using Table 2, calculate the following:

a. What is the impact of the program?

b. Is this difference in outcomes an accurate estimate of the impact of the program? Why or why not?

2.3 Suppose we have a strong reason to believe that the true treatment effect on employment is large, positive, and significant. How might the following scenarios influence our ability to accurately estimate the treatment effect?

a. The 20% of the sample with the lowest income across both the comparison and treatment groups leave the study.

b. In the comparison group, the individuals with the most job experience are able to find work abroad and migrate, exiting the study.

DISCUSSION TOPIC 3: PARTIAL COMPLIANCE

In the JeA study, random assignment determined who among the eligible applicants would be offered a spot in the vocational training program (and likewise, who would not be offered a spot). However, not everyone in the treatment group followed through by attending the program, and some people in the comparison group managed to attend the program even though they were not formally admitted. In research parlance, these two groups of people are called "non-compliers" because they do not comply with their treatment assignment, and we would say that we have partial compliance in the sample overall. In the study, the rate of non-compliance was low, but for the purposes of this case study we will consider scenarios with higher rates of non-compliance.[3] In this section, we will examine the consequences of partial compliance and how to prevent or minimize this threat.

TABLE 3A: EARNINGS POST INTERVENTION BY TREATMENT ASSIGNMENT

Assignment	Average earnings	Number of individuals
Treatment	$295,300	4,000
Comparison	$260,000	2,000

3.1 Imagine you compare the earnings of those *assigned to* the treatment group to the earnings of those *assigned to* the comparison group, regardless of whether they comply with their treatment assignment. What is the impact of the treatment?

[3] The researchers and ECAPs worked in tandem to ensure only those who were assigned to treatment were able to attend the JeA courses and that those selected to enter the program had a high likelihood of completing the program, resulting in a high rate of compliance. In total, the rate of compliance was 97%, meaning only 3% of the sample either got the treatment when assigned to the comparison group OR did not take up the treatment when assigned to the treatment group. Since non-compliance was low, it isn't a concern for this particular study but is a topic still worth exploring.

3.2 Using administrative data from their ECAP partners, the researchers could identify which individuals actually attended training courses, regardless of their treatment assignment:

TABLE 3B: TREATMENT ASSIGNMENT VS. TREATMENT STATUS			
	Assignment		
Take-up status	Treatment	Comparison	Total
Attended the JeA program	3,500	100	3,600
Did not attend the JeA program	500	1,900	2,400
Total	4,000	2,000	6,000

a. Some of your colleagues are passing by your desk and say you should consider comparing the 3,600 individuals who attended the program to the 2,400 individuals who did not attend the JeA program. Is this advice sound? Why or why not?

b. Other colleagues believe that you should calculate the effect of the treatment by comparing the 3,500 people who were assigned to and attended the JeA program to the 1,900 people who were not assigned to and did not attend the program. Is this advice sound? Why or why not?

c. Another colleague suggests that you use the compliance rates, the proportion of people in each group that did or did not comply with their treatment assignment. You should divide the "intention to treat" estimate by the difference in treatment ratios (i.e., proportions of each experimental group that received the treatment). Is this advice sound? Why or why not?

3.3 Using this new information on compliance, calculate the proportion of the treatment group who were treated and the proportion of the comparison group who were treated. Use your estimate of the ITT from question 1 and the information in the table to estimate the LATE, as follows:

$$LATE = \frac{ITT}{proportion\ of\ treatment\ group\ who\ took\ up\ treatment - proportion\ of\ comparison\ group\ who\ took\ up\ treatment}$$

3.4 Is the LATE bigger or smaller than the ITT? Why would the LATE be different from the ITT?

APPLICATIONS TO OTHER CONTEXTS

While we focused on the Jóvenes en Acción example in this case study, the evaluation design, intervention, and its findings have relevance to other contexts as well. The findings from this evaluation show that vocational training can have large, lasting, and cost-effective results—and that administrative data can be key to measuring such effects.

Many countries face the issue of high youth unemployment rates. In Egypt, following the Arab Spring, rates were as high as 35% (Elsayed et al. 2018); in the United States, since 2012, rates have hovered between 9.6% to 18% (Davis and Heller 2020); in France as of 2015, the youth unemployment rate was double the national average (Crépon et al. 2015). Attanasio et al. showed the three-part JeA approach (skills training, internships, and application coaching) was an effective way to reduce youth unemployment both in the short-term (2011) and long-term (2017). Governments looking to reduce youth unemployment may want to consider a similar model.

One reason Attanasio and co-authors were able to show such long-term effects was by partnering with the Colombian government to gain access to administrative data. This tax data facilitated a long-term follow up on a larger sample at a lower cost. Governments and other organizations looking to evaluate their programs may want to consider how they can use their administrative data to facilitate evaluations.

REFERENCES AND FURTHER READING

Attanasio, Orazio, Adriana Kugler, and Costas Meghir. 2011. "Subsidizing Vocational Training for Disadvantaged Youth in Colombia: Evidence from a Randomized Trial." *American Economic Journal: Applied Economics* 3 (3): 188–220. DOI: 10.1257/app.3.3.188

Attanasio, Orazio, Arlen Guarin, Carlos Medina, and Costas Meghir. 2017. "Vocational Training for Disadvantaged Youth in Colombia: A Long-Term Follow-up." *American Economic Journal: Applied Economics* 9(2): 131-143. DOI: 10.1257/app.20150554

Crépon, Bruno, Marc Gurgand, Philippe Zamora, Roland Rathelot, and Thomas Le Barbanchon. 2015. "Accès a` l'apperntissage et sécurisation des parcous: bilan d'une expérimentation aléatoire contrôlée conduit sur 7 sites en France." Rapport d'Evaluation CREST-DARES CREDOC.

Davis, Jonathan M.V., and Sara B. Heller. 2020. "Rethinking the Benefits of Youth Employment Programs: The Heterogeneous Effects of Summer Jobs." *The Review of Economics and Statistics* 102(4): 664-677. DOI: 10.1162/rest_a_00850

Elsayed, Ahmed, Kevin Hempel, and Adam Osman. 2018. "Overcoming Youth Unemployment in Egypt: Randomized Evaluations Showcase the Promise of Active Labor Market Programs." Working Paper.

EGAP. 2020. "10 Things to Know About the Local Average Treatment Effect." Last accessed May 23, 2024.

J-PAL. "Data Analysis." J-PAL Research Resource. Abdul Latif Jameel Poverty Action Lab.

UNESCWA. 2022. "Statistical Terms Glossary." Last accessed May 23, 2024.

REUSE AND CITATIONS

To request permission to reuse this case study or access the accompanying teachers' guide, please email training@povertyactionlab.org. Please do not reuse without permission. To reference this case study, please cite as:

> J-PAL. 2024. "Case Study: Vocational Training in Colombia: Threats and Analysis." Abdul Latif Jameel Poverty Action Lab.

To reference the original studies by Attanasio and coauthors, please cite as:

> Attanasio, Orazio, Arlen Guarin, Carlos Medina, and Costas Meghir. 2017. "Vocational Training for Disadvantaged Youth in Colombia: A Long-Term Follow-up." *American Economic Journal: Applied Economics* 9(2): 131-143. DOI: 10.1257/app.20150554.

Attanasio, Orazio, Adriana Kugler, and Costas Meghir. 2011. "Subsidizing Vocational Training for Disadvantaged Youth in Colombia: Evidence from a Randomized Trial." *American Economic Journal: Applied Economic*s, 3 (3): 188-220. DOI: 10.1257/app.3.3.188