Dynamic Treatment Effects of Job Training*

Jorge Rodríguez[1], Fernando Saltiel[2], and Sergio Urzúa[3]

[1]Universidad de los Andes
[2]Duke University
[3]University of Maryland and NBER

November 8, 2019

Abstract

This paper estimates the dynamic returns to job training. We posit a model of sequential training participation, where decisions and outcomes depend on observed and unobserved characteristics. We analyze different treatment effects, including policy relevant parameters, and link them to continuation values and latent skills. The empirical analysis exploits administrative data combining job training records, matched employee-employer information, and pre-labor market ability measures from Chile. Although the average returns to training are small, these vary across the unobserved ability distribution and previous training choices. In fact, among young workers, the returns to training are lower when followed by additional training. Thus, we provide evidence of dynamic substitutability and examine potential mechanisms driving this result. For instance, policy experiments illustrate how increasing the local availability of training programs may affect earnings heterogeneously across dynamic responses.

*Jorge Rodríguez, Universidad de los Andes, Chile; email: jrodriguezo@uandes.cl; Fernando Saltiel, Department of Economics, University of Maryland; email: fernando.saltiel@duke.edu; Sergio Urzúa, Department of Economics, University of Maryland and NBER; email: urzua@econ.umd.edu. We thank Matias Cattaneo, Xavier D'Haultfoeuille, Claudio Ferraz, Maurice Kugler, Thomas Le Barbanchon, Lance Lochner, Matt Masten, Arnaud Maurel, Oscar Mitnik and Gabriel Ulyssea for useful comments and suggestions. We are indebted to seminar participants at IZA/CREST/OECD Conference on Labor Market Policy Evaluation (2018), George Mason University (2018), Latin American Meeting of the Econometric Society (2017), RIDGE Impact Evaluation of Labor Market Policies Conference (2017), PUC-Chile (2017), *Banco de Desarrollo de America Latina* (2016), PUC-Rio (2016), and LACEA-LAMES (2016). We thank the Ministry of Finance of Chile for providing us access to the data used in this paper.

1 Introduction

The evolving complexity and uncertainty in the demand for skills in a changing labor market induces workers to constantly revise their human capital investment decisions. Recent research has documented workers may possess different sets of competencies vis-á-vis those required in the workplace.[1] In this context, understanding the dynamics of training decisions and their associated returns has gained prominence over the past years: as fast technological progress shifts the set of skills that jobs demand, workers may participate in on- or off-the-job training multiple times in their careers.

This paper estimates the returns to job training in a context of dynamic training choices and labor market outcomes. We use our framework to provide new insights into the static and dynamic effects of job training on earnings, including the identification of heterogeneous responses across different groups as well as an empirical assessment of continuation values and dynamic complementary (substitutability) arising from repeated training participation.[2] Ultimately, we not only document who benefits from training but also how the timing and sequence of the training decisions define potential and actual gains (and losses). Furthermore, we present and estimate dynamic policy relevant treatment effects, which consider how policy changes in one period affect long-term earnings by potentially altering workers' present and future training decisions.

In the spirit of Heckman and Navarro (2007) and Heckman et al. (2016), we posit a tractable dynamic-discrete choice model of job training. In the model, a worker must decide whether or not to take part in a training course across multiple periods, and workers may participate in training on multiple occasions. For a given training history, and conditional on firm characteristics, the agent chooses to participate in a job training course if the perceived net benefits are positive. Individual choices and outcomes depend on observed characteristics as well as on unobserved heterogeneity, which we interpret as the initial stock of skills. Using a measurement system of pre-labor market test scores, we are able to nonparametrically identify the distribution of latent skills and use it to identify the joint distribution of counterfactual earnings across potential training choices.

[1] Guvenen et al. (2015), Postel-Vinay and Lise (2015), Lise et al. (2016), and Saltiel et al. (2018) discuss skill (mis)match in the labor market. See Autor et al. (2003), SpitzOener (2006), Ingram and Neumann (2006), Acemoglu and Autor (2011), and Sanders and Taber (2012) for a literature on the changing returns to specific skills.

[2] Formally, if we let Y_t be earnings at the end of period t, the production function of Y_t exhibits dynamic complementarity (substitutability) if the return to D_t (training participation at time t) is higher (lower) conditional on prior training participation, i.e. $\frac{\partial^2 Y_t}{\partial D_{t-1} \partial D_t} \geq (\leq) 0$.

We exploit a large-scale training program in Chile called "*Franquicia Tributaria*" (FT). FT fully subsidizes training courses at off-site providers for workers who are employed in a formal-sector firm. In the program, a worker can participate in a training course on multiple occasions, and almost half of workers do so. We take advantage of administrative data on job training records for the population of labor market entrants from years 2003-2008 and combine it with matched employee-employer data on labor income. We augment these data with measures on workers' pre-labor market abilities coming from college admission test scores. Since workers with no prior job training experience may find it especially valuable to take up job training, as they anticipate higher returns to human capital investments, we analyze the earnings returns to training for first-time labor market entrants.

Using the estimated model parameters, we document static and dynamic treatment effects of job training. We first examine the effects of training on workers' earnings in the first two years in the labor force. The static returns to training indicate that program participation in the first year raises average monthly earnings by 1.7%. Meanwhile, the returns to second-period participation differ across first-period histories, reaching 3.4% for non-trained workers while remaining below 1% for first-period trainees. To consider the dynamic returns to training, we assess the impact of first-period training on present discounted value of earnings across two years. The dynamic average treatment effect (DATE) indicates that first-period training increases the present value of earnings by 3.6%. We examine the mechanisms through which early training delivers positive medium-term impacts by decomposing the DATE parameter into the direct effect of training and its continuation value, which links human capital investment decisions and potential gains over time. We find that while the short- and medium-term direct effects are positive and significant, the continuation value is not statistically different from zero. Furthermore, we find evidence of dynamic substitutability, as training in the first period reduces the earnings returns to training in the second period. Furthermore, dynamic substitutability is stronger for high-skilled workers. We note that dynamic substitutability may be explained by the structure of the job training courses examined in this paper, or more generally by the production function of post-schooling human capital accumulation. These results indicate that the post-school human capital production function differs from that of school-age children, which instead exhibits dynamic complementarities (Cunha et al., 2006; Johnson and Jackson, 2017).

To examine the policy implications arising from these results, we estimate the effect of an increase in local course-hour availability. As this policy may affect workers' choices in both years, we identify dynamic response types, defined by the reactions to the policy change in each time period. For instance, we can identify a group of workers who are induced to participate in training in both time periods due to the policy change, despite being baseline never-participants. We find that a 10% increase in course availability would increase the medium-term earnings of affected workers by 3%, with largely homogeneous effects across the skill distribution. Moreover, we find similar-sized impacts for larger program expansions, yet document that the effects are heterogeneous across dynamic response types, as workers induced to participate in both periods would enjoy the largest gains from the policy change.

Our paper contributes to an extensive literature analyzing the effect of job training programs on labor market outcomes. In a non-experimental context, the inherent identification challenge arises from potential self-selection into training. To deal with this concern, various papers have relied on individual fixed effect estimators to account for unobserved individual heterogeneity (Ashenfelter, 1978; Ashenfelter and Card, 1985; Heckman and Hotz, 1989; Lynch, 1992; Booth, 1993; Veum, 1997; Lengermann, 1999; Frazis and Loewenstein, 2007; Mueser et al., 2007; Albert et al., 2010). While Heckman et al. (1998, 1999) show that the standard fixed-effect estimator can effectively remove selection bias, this estimation strategy does not account for dynamic selection into training or estimate heterogeneous returns (Callaway and Sant'Anna, 2018; de Chaisemartin and D'Haultfœuille, 2018; Goodman-Bacon, 2018). A parallel strand of the literature has estimated the effects of training using matching estimators (Heckman et al., 1997; Smith and Todd, 2005; Andersson et al., 2013; Lechner, 2000; Larsson, 2003; Dyke et al., 2006; Lechner and Wunsch, 2009). Our empirical strategy extends this analysis by allowing for matching on unobserved characteristics, while incorporating exclusion restrictions in the training participation decision. Furthermore, we estimate heterogeneous returns to training across workers' latent ability and allow for the effects to vary across training histories and time periods, extending in this way the existing job training literature.

This paper also contributes to a growing literature on dynamic treatment effects. Previous studies have estimated the dynamic returns to educational attainment, including Heckman and Navarro (2007), Heckman et al. (2016), and Heckman et al. (2018), among others. To the best of our knowledge, this is the first paper to estimate dynamic returns to post-schooling human capital accumulation, in the context of job training for employed workers.[3] Moreover, we present the first estimates of the continuation value and dynamic complementarities of job training, allowing us to explore how early training affects the returns to additional training stints. By showing that job training participation exhibits dynamic substitutability, we present a novel difference relative to the existing evidence on human capital accumulation during formal schooling. Lastly, we contribute to a growing literature exploring policy relevant treatment effects by defining dynamic response types and estimating the heterogeneous impacts of a policy change across groups (Heckman and Vytlacil, 2001, 2005; Carneiro et al., 2010; Mogstad et al., 2018; Mogstad and Torgovitsky, 2018). We present the first estimates of how increased early-career job training availability may affect earnings by shifting workers' subsequent training participation.

The rest of the paper is organized as follows. Section 2 presents a Dynamic Roy model and discusses the relevant treatment effects. Section 3 describes the institutional setup, data sources, sample characteristics and presents reduced-form estimates. Section 4 specifies our model of sequential training participation with unobserved heterogeneity, presents estimated model parameters, considers goodness of fit, and documents the implied patterns of selection on unobservables. Section 5 defines the static and dynamic treatment parameters and presents evidence on the returns to job training. Section 6 discusses the simulated policy intervention, defines and identifies dynamic response types and presents evidence on the returns to the policy change. We conclude in Section 7.

2 Treatment Effect Framework

In this section, we use a Roy model framework to characterize the dynamics of training decisions and labor market outcomes. The model considers training decisions across multiple periods and allows earnings counterfactuals to vary freely across all potential histories of training choices. Within this framework we define the treatment effects of interest, link them to continuation values, and discuss dynamic complementarity/substitutability of training investments.

[3] A parallel strand of the literature has estimated dynamic treatment effects in the context of job training for the unemployed, making these papers different in scope from our analysis (Abbring and van den Berg, 2003; Fitzenberger and Völter, 2007; Fredriksson and Johansson, 2008; Fitzenberger et al., 2016; Ba et al., 2017).

2.1 Dynamic Roy Model

The essence of the model involves an agent making training choices for many periods and earnings, which directly depend on previous decisions. In any period t, potential earnings depend on her current training decision as well as on the entire history of training activities.[4] In period t, the agent makes her optimal training decision, and she is allowed to participate in job training as frequently as desired.

We model the dynamic training decision as a tree of sequential binary decisions, where the individual chooses training in each stage $t \in \mathcal{T} \equiv \{1, ..., T\}$. We define $\mathcal{H}_t$ as the set of all possible training decisions histories up through time t. An element in that set, $h_t \in \mathcal{H}_t$, represents a specific training history, not including the training decision to be taken in period t. Thus, h_1 is the initial condition. As workers have not been able to take up training prior to entering the labor force in the first period, h_1 can be characterized by an empty set. To illustrate this notation, Figure 1 depicts the decision tree for $T = 2$. In period $t = 2$, $\mathcal{H}_2 = \{0, 1\}$, where $h_2 = 1$ if the agent was trained in $t = 1$ and 0 otherwise. Likewise, $\mathcal{H}_3 = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$, where each element $h_3 \equiv (l, j)$ denotes the training decision in $t = 1$ and $t = 2$, respectively.

At each choice node, agent i compares the benefits and costs of the available alternatives to make her next training choice. Let $D_i(h_t)$ be her training decision in period t given history $h_t \in \mathcal{H}_t$. $D_i(h_t)$ equals one if she decides to participate in a training program, and zero otherwise. Her optimal choice is given by:

$$D_i(h_t) = \begin{cases} 1 & \text{if } I_i(h_t) \geq 0 \\ 0 & \text{otherwise} \end{cases} \qquad h_t \in \mathcal{H}_t,\ t \in \mathcal{T} \tag{1}$$

where $I_i(h_t)$ denotes the value of training in period t for a given history $h_t \in \mathcal{H}_t$. $I_i(h_t)$ may incorporate non-pecuniary benefits and costs of training. In principle, expression (1) provides a general framework. It can accommodate, for example, forward-looking agents anticipating and acting based on present and future benefits of training in period t, who are uncertain with respect to the true model that generates counterfactual earnings. The agent progresses through each node

[4] We focus on training choices for workers who do not fall into unemployment. We thus abstract from modeling the employment decision, which could affect the probability of training. We note the model could generally accommodate unemployment as an outcome variable, however.

after making training choices, and for each possible choice and training history, there is an associated labor market outcome. Let $Y_i(h_t; j)$ be potential earnings for a training decision $j \in \{0, 1\}$ made by worker i with history h_t. As is common in the literature, it captures workers' earnings immediately after job training participation.

2.2 Returns to Job Training

For workers, effective job training participation can have immediate effects on earnings. Using our notation, we define this (static) individual-level impact as $Y_i(h_t; 1) - Y_i(h_t; 0)$. However, job training may also affect workers' long-term earnings, both directly, as trained workers have potentially accumulated human capital which increases their labor market productivity, but also indirectly, by shifting workers' training participation in subsequent periods. To investigate these channels, we consider the earnings stream of individual i facing the training decision in period t. Formally, let $\widetilde{Y}_i(h_t; j)$ be her present value of earnings associated with training option j given history h_t. Thus,

$$\widetilde{Y}_i(h_t; j) \equiv Y_i(h_t; j) + \rho \left[\widetilde{Y}_i(h_{t+1}; 0) + D_i(h_{t+1}) \left[\widetilde{Y}_i(h_{t+1}; 1) - \widetilde{Y}_i(h_{t+1}; 0) \right] \right], \qquad j \in \{0, 1\},\ h_t \in \mathcal{H}_t$$

ρ is a discount factor, $D_i(h_{t+1})$ takes a value of one if worker i participated in training at $t + 1$ and $\widetilde{Y}_i(k; h_{t+1})$ captures the earnings stream associated with training decision k at $t + 1$ given history h_{t+1}. Thus, individual-level dynamic treatment effect of participating in training in period t can be expressed as $\widetilde{Y}_i(h_t; 1) - \widetilde{Y}_i(h_t; 0)$. To understand the different mechanisms through which job training affects long-term earnings, we follow Heckman et al. (2016) and decompose it as:

$$\begin{aligned}\widetilde{Y}_i(h_t; 1) - \widetilde{Y}_i(h_t; 0) = & \underbrace{(Y_i(h_t; 1) - Y_i(h_t; 0))}_{\text{Direct effect (short-term)}} + \underbrace{\rho[\widetilde{Y}_i(h'_{t+1}; 0) - \widetilde{Y}_i(h_{t+1}; 0)]}_{\text{Direct effect (long-term)}} + \\ & \underbrace{\rho \left[D_i(h'_{t+1})(\widetilde{Y}_i(h'_{t+1}; 1) - \widetilde{Y}_i(h'_{t+1}; 0)) - D_i(h_{t+1})(\widetilde{Y}_i(h_{t+1}; 1) - \widetilde{Y}_i(h_{t+1}; 0)) \right]}_{\text{Continuation value}},\end{aligned} \tag{2}$$

where the sequences of decisions contained in h'_{t+1} and h_{t+1} differ only in the training decision observed in period t ($h'_{t+1} = (h_t, 1)$ and $h_{t+1} = (h_t, 0)$). The first two terms of the right-hand side capture the direct effect of training at t (properly discounted). The first term recovers the direct effect of training on earnings immediately following participation. The second term represents the impact of training at time t on lifetime earnings conditional on not taking up training at $t+1$. This parameter recovers the direct effect of the baseline training stint without considering additional gains arising from future training participation. The third term, which recovers the additional gain (if any) of training in $t+1$ from training in t for those who take up training in $t+1$, corresponds to the continuation value of job training. As illustrated in Section 4, continuation values can be defined in settings with at least two sequential training decisions and a resulting outcome.

Since continuation values are informative of how job training may result in an increase in long-term earnings, they been previously estimated in the human capital investment literature (Heckman et al., 2018). However, for an econometrician interested in understanding whether job training leads to larger/smaller earnings gains arising from dynamic complementarity/substitutability, the continuation value will not directly recover this parameter. To see this, consider the following decomposition:

$$\begin{aligned}
\underbrace{D_i(h'_{t+1})(\widetilde{Y}_i(h'_{t+1};1)-\widetilde{Y}_i(h'_{t+1};0))-D_i(h_{t+1})(\widetilde{Y}_i(h_{t+1};1)-\widetilde{Y}_i(h_{t+1};0))}_{\text{Continuation Value}} &= \\
\underbrace{(\widetilde{Y}_i(h'_{t+1};1)-\widetilde{Y}_i(h'_{t+1};0))-(\widetilde{Y}_i(h_{t+1};1)-\widetilde{Y}_i(h_{t+1};0))}_{\text{Dynamic Complementarity/Substitutability}} &+ \\
\underbrace{(D_i(h'_{t+1})-1)(\widetilde{Y}_i(h'_{t+1};1)-\widetilde{Y}_i(h'_{t+1};0))-(D_i(h_{t+1})-1)(\widetilde{Y}_i(h_{t+1};1)-\widetilde{Y}_i(h_{t+1};0))}_{\text{Dynamic Sorting Gains}} &. \qquad (3)
\end{aligned}$$

Thus, the continuation value of training equals dynamic complementarity/substitutability plus a sorting term we label "dynamic sorting gains." As noted above, dynamic complementarity (substitutability) is informative about the production function of human capital of training across multiple periods, exhibiting dynamic complementarity (substitutability) if the return from training in, for example, $t+1$ is higher (lower) conditional on time t participation. Therefore, when the continuation value is larger than the dynamic complementarity (substitutability), workers are positively sorting into training participation (positive dynamic sorting gains).

We have so far discussed individual-level parameters which are informative about the static and dynamic returns to job training participation. Recovering them is a difficult endeavor as individuals can endogenously sort into training based on (observed and unobserved) potential benefits, and may do so across multiple time periods. We further note that both training participation and the returns to training may directly depend on workers' underlying baseline productivity. For instance, less productive workers may be more likely to take up job training, while enjoying the largest returns from program participation. As a result, any empirical strategy must account for endogenous dynamic program participation, consider the heterogeneous returns to job training, and estimate the different parameters introduced in this section to correctly capture the various benefits arising from job training. In Section 4, we introduce a dynamic discrete choice model of job training participation decisions. In this model, we proxy for baseline productivity using measures of pre job training ability, allowing us to estimate heterogeneous static and dynamic returns to training as well as policy-relevant treatment effects. In the next section, we present our data sources, introduce the training program under consideration and examine whether reduced-form strategies can recover static and dynamic returns to training.

3 Data Sources and Descriptive Analysis

3.1 Institutional Context

In this paper, we examine the effects of a nationwide funding scheme for job training programs called *Franquicia Tributaria* (FT) in Chile. The program funds training courses for a significant number of workers in the country in any given year through a large-scale subsidy for training expenditures undertaken by firms.[5] As a result, the program targets formal-sector workers. While all workers are theoretically eligible for the program, the design of the funding scheme implies that those at larger firms are more likely to participate in training through FT.[6] Training courses are held off-the-job, in centers managed by private providers.[7] There are three types of training courses covered by FT: (i) short-term courses, including industry-specific programs (such as learning to operate heavy machinery), general-skills courses (such as Microsoft Excel courses), as well as programs focused on

[5] In 2010, FT funded training courses for 920,688 workers, covering 12% of the labor force. Using data from Chile's national accounts and from Carrillo et al. (2018), we estimate the total cost of FT to be 0.08% of GDP.

[6] FT subsidizes firms through a tax exemption, with a cap set at 1% of the firms' annual payroll. Carrillo et al. (2018) show that the government subsidizes 80% of the total cost of courses, making FT courses relatively inexpensive for firms. The structure of the FT subsidy implies that the cost to medium- and large-sized firms is significantly smaller than for firms with less than ten employees. Since our model examines the extent to which workers self-select into training, we restrict our attention to workers employed at firms facing negligible costs for FT courses.

[7] There are over 15,000 providers in the market. Courses are generally scheduled after-work and on weekends, ensuring they do not interfere with regular work schedules.

soft skills; (ii) short-term degrees leading to specialization in specific disciplines; (iii) professional conferences.[8] The predominant role played by the private sector in training course provision under FT resembles that of the Workforce Investment Act (WIA) in the United States.[9] While the workers targeted by the FT program are different than those in the WIA, we note that our analysis speaks to the effectiveness of short-term courses that are also commonly used in other contexts.

3.2 Data Sources

To recover training histories and associated labor market outcomes, we construct a novel database that merges three different sources of information. First, we take advantage of administrative records from *Franquicia Tributaria*. Using this data source, we construct workers' training histories by observing their participation in FT-subsidized courses from 1998 through 2010. We analyze labor market outcomes using information from Chile's Unemployment Insurance (UI) system. UI data registers workers' monthly earnings and the firm of employment for all workers with formal sector contracts. We focus on the worker's main employment stint in each quarter and examine earnings in the first quarter of the year following the training event. For notational simplicity, we refer to the outcome variable as contemporaneous with the training event at time t. Lastly, our final source of information comes from performance in a college-entry examination (PSU), which is a mandatory test for all students who wish to enter a post-secondary institution. We observe PSU scores for all high school graduates who took the test between 2000 and 2007. Using individual identifiers, we recover PSU scores of workers to supplement our data of labor market outcomes and training choices. We work with standardized PSU test scores (computed separately by year). The PSU database also includes information on student's observable characteristics, such as gender, age, parental education, family size, and parental employment at the time of the test.

To circumvent threats to identification and for computational tractability, we restrict our sample in several ways. First, we focus our attention on the returns to multiple job training courses for young workers who are first-time labor entrants. We impose this restriction given that we do not

[8] Courses of type (i) and (iii) must be of five hours at least while courses type (ii) must be over 100 hours. In 2009, the duration of the average job training course was 19.3 hours (Comisión Revisora del Sistema de Capacitación e Intermediación Laboral, 2011).

[9] One of the main goals of the WIA (which replaced the Job Training Partnership Act), was to strengthen the role of the private sector (Barnow and Smith, 2016). In practice, the WIA defined individual training accounts (vouchers for training) by which individuals can choose short-term, off-the-job, training courses held in private providers (Andersson et al., 2013).

observe training histories before 1998; if training choices depend on prior training decisions, then we would be omitting a relevant variable—past training—in the choice equation.[10] We use the sample of young workers, identify their first year of employment, and follow their labor market history thereafter—by definition, their job training history in our first period is zero. Second, for tractability, we restrict our analysis to training stints during their first two years in the labor force and examine extensive-margin training decisions on a yearly basis. As a result, workers are trained at most twice during our period of interest. Third, we restrict the sample to individuals who are eligible to participate in training financed by FT—that is, individuals who work in the formal sector. As our analysis of worker self-selection into training requires workers to be able to take part in courses each year if they want to, we limit the sample to individuals who are employed for at least nine months in each of their first two years in the labor force in firms with at least ten employees. Since UI data indicates that 90% of formal-sector employment in Chile is at firms with at least ten employees, this restriction is not necessarily binding. In this way, we analyze a group of workers who are effectively eligible for training each year. By doing so, we abstract from analyzing effects of training on employment.[11] We focus on labor market entrants from years 2003-2008 and their training choices two years after entry. Our final sample consists of 37,089 workers who meet all of the above criteria.

3.3 Descriptive Statistics

Table 1 presents summary statistics for our sample. 54% of individuals in our sample are women. The age at the time of taking the college entrance exam is in line with the average for the country (close to 18 years old). The average GPA and PSU scores in math and language are largely in line with the national average. The average monthly salary in the first quarter of the second year in the labor force equals 551 dollars, reaching 624 dollars after the second year in the labor market.

Table 2 reports summary statistics across workers' training participation in their first two years in the labor force. A training group is denoted as (h, h'), where the first and second entries denote training participation in the first and second period, respectively $(h, h' \in \{0, 1\})$. The never-trained group is by far the largest in our sample, representing 61 % of all individuals. This group has a lower

[10] This omission could influence the estimation of the distribution of unobserved heterogeneity. See Heckman (1981) and Arellano and Honore (2001) for a related discussion.

[11] Our focus on the earnings dimension makes our paper comparable in scope to the on-the-job training literature.

PSU and GPA than workers in all other groups, with the largest difference appearing relative to always-trained workers. The unconditional earnings differential between these two groups reaches 48 and 45% for the first and second period, respectively. We highlight earnings differentials across workers trained only in the first year relative to second year trainees, as in a world of constant returns to training, earnings differentials should not appear for these workers. In this context, later-trained workers earn higher salaries than early-trainees, despite similar test score performance. These intertemporal earnings differences across groups with similar stock of training suggest the presence of different treatment effects of job training over time.

3.4 Reduced-Form Analysis

The descriptive statistics presented above highlight important differences in baseline characteristics and outcomes across training participation groups. We further explore whether training is associated with increased earnings by first estimating OLS regressions of the short-term returns to training. We present the results in the first panel of Table 3. In the first two columns, we examine the returns to training in the first year in the labor force on earnings in the first quarter of the following year (defined as Y_{i1} in Section 2). OLS regressions suggest that training participation increases earnings by 15-18%, with lower estimated impacts upon controlling for test scores and background characteristics. We find different effects in the regressions of second-period training on earnings (presented in columns 3 and 5) as this event increases earnings by 25.7% and 20.9% conditional on having trained and not trained in the first period, respectively. Nonetheless, upon the inclusion of control variables, we find similar short-term returns to training, in the 15-18% range.

The OLS estimates presented in Table 3 can only be interpreted as causal based on the strong assumption of selection on observables. Since this assumption is unlikely to hold, the job training literature (summarized in Card et al. (2010)) has previously taken advantage of the longitudinal aspect of the data to estimate the returns to training in the following equation:

$$Y_{it} = \delta D_{it} + X_{it}'\beta + \alpha_t + \kappa_i + \varepsilon_{it}, \tag{4}$$

where Y_{it} represents the log of earnings for worker i in period t (year since labor market entry).

D_{it} is a dummy that equals 1 if individual i was trained at time t and zero otherwise, α_t is a year fixed effect and κ_i an individual fixed effect. δ captures the average effect of training on earnings in period t, yet it does not recover the direct impacts of any particular training event.

Panel B in Table 3 presents the estimated results following from different versions of equation (4). Column 1 shows that concurrent training participation is associated with an earnings gain in the range of 22%. In the second column, we add control variables and find the returns to training remain both economically and statistically significant, reaching 18% — similar in magnitude to the estimates presented for each training event in Panel A. In the third column, we add an individual fixed effect. We find that the estimated impact of job training is significantly lower, falling to 0.9%, though remaining statistically significant. This result fits in with previous findings by Frazis and Loewenstein (2007), who show that controlling for unobserved heterogeneity through an individual fixed effect largely attenuates the estimated impacts of training in an OLS regression. In the last column, we include firm fixed effects to explore whether the returns to training are explained by workers sorting into higher-quality firms.[12] We do not find major changes in the estimated returns to training.

While the fixed-effect estimation in Table 3 may eliminate selection bias, this regression might not identify a pre-determined parameter of interest. First, Goodman-Bacon (2018) shows that in setting where the timing of treatment varies, the usual fixed effect estimator recovers a weighted average of all possible pairs of the underlying differences-in-differences estimator. Moreover, when treatment effects are not constant, some of these weights might be even negative (de Chaisemartin and D'Haultfœuille, 2018). In the context of our model of choices and counterfactual outcomes, Web Appendix A defines and tests the assumptions needed for fixed-effect estimators to recover the average treatment effect of training. A first-difference estimator recovers the ATE of training only if: (i) the earnings returns from training are constant across time and training histories and, (ii) the earnings returns from training do not vary with unobserved ability. In Web Appendix A, using our model estimates, we find evidence against the null hypothesis of lack of differential gains at conventional significance levels.[13] Second, the parallel-trends assumption required for identification

[12] To examine the importance of individual unobserved heterogeneity and firm quality to the change in the estimated returns to training from Column 2 to Column 4, we estimate a decomposition following Gelbach (2016). We find that firm fixed effects account for less than 5% of the change in the estimated returns to training from column (2) to (4), with the vast majority being explained by individual fixed effects.

[13] We present reduced-form evidence of heterogeneous effects of training by estimating an OLS regression of different

in reduced-form analyses may not hold if agents sort into training based on unobserved gains from training. Third, even under a parallel trends assumption and assuming a correctly re-weighted fixed-effect estimator (Callaway and Sant'Anna, 2018), this estimator would not recover direct effect, continuation value and dynamic complementarity/substitutability presented in equations (2) and (3). As a result we would still miss potentially important questions for understanding the nature of job training programs: To what extent do workers self-select into training based on unobserved characteristics? What are the dynamic returns to job training? Are there heterogeneous returns to training across workers' latent ability? What is the role of continuation values and dynamic complementarity (substitutability) in the returns to training? To this end, we present our dynamic discrete choice model in Section 4.

4 Empirical Implementation of Dynamic Model

This section discusses the implementation of the model and presents estimated model parameters. We first discuss the model specification, its implementation and the identification of the distribution of unobserved ability. We adopt the analysis of Heckman and Navarro (2007) and Heckman et al. (2016) to the context of job training. As discussed below, we identify the model non-parametrically through the combination of a matching-on-unobservables assumption and exclusion restrictions. We then present the estimated parameters and analyze the extent of sorting on unobserved ability.

4.1 Model Specification and Identification

As discussed in Section 2, individuals decide to participate in job training based on the net value of training in period t for a given training history $h_t \in \mathcal{H}_t$, defined by $I_{it}(h_t)$. We assume that the value of training depends on individuals' observed and unobserved characteristics, and we model the decision process using a linear-in-the-parameters equation in:

$$I_i(h_t) = X_i^I \beta^I(h_t) + \eta_i^I(h_t) \qquad h_t \in \mathcal{H}_t,\ t \in \mathcal{T}, \tag{5}$$

training histories against quarterly earnings following the second year in the labor force (Table A.2). Confirming the results in Table 2, always trained students outearn both their never-trained and once-trained peers—yet the main difference appears among workers who are trained only once: late-trainees outearn first-year training participants by 8.1%. This analysis is similar to the estimation of multi-valued treatment effects, as in Cattaneo (2010).

where X_i^I is a vector of individual characteristics (observed by the econometrician and the agent) and $\eta_i^I(h_t)$ is an individual- and choice-level innovation that the agent uses to make current choices but it is otherwise unobserved by the analyst. We do not explicitly specify individual preferences and expectations formation. Hence, forward-looking behavior is not imposed in the model. Since we do not model preferences and budget sets, we abstract away from assumptions about behavior and uncertainty—ubiquitous elements in the structural literature (Keane et al., 2011).

The potential outcome associated with $j \in \{0,1\}$ and $h_t \in \mathcal{H}_t$ also depends on observed and unobserved characteristics, and it is given by:

$$Y_i(h_t;j) = X_i^Y \beta^Y(h_t;j) + \eta_i^Y(h_t;j) \qquad j \in \{0,1\}, h_t \in \mathcal{H}_t,\ t \in \mathcal{T}, \tag{6}$$

where X_i^Y is a vector of observed characteristics and $\eta_i^Y(h_t;j)$ reflects latent productivity, which is unobserved by the econometrician.

Note that the effects of observed characteristics and unobserved ability on the value of training and outcomes vary across time periods and training histories. Both the parameters of the training choice and earnings equation are allowed to vary with training counterfactual choices. Thus, our model extends the conventional setting (where past outcomes and choices sometimes enter as lagged variables): choices depend on past decisions and and earnings vary accordingly. Furthermore, the earnings return to different paths of training choices—defined by h_t and j—-vary across the latent ability distribution.

Heckman and Navarro (2007) provide the basis for identification in a general dynamic discrete-choice model, an argument that can applied in our context.[14] Both exclusion restrictions and matching on unobserved heterogeneity allows for identification of the joint distribution of outcomes. This result enables us estimating parameters such as ATE, TT, and TUT, both in static and dynamic contexts. For sake of brevity, we do not reproduce and prove the main identification theorems from Heckman and Navarro (2007). Instead, we discuss the intuition and necessary conditions applied to our context.

By applying Theorem 2 of Heckman and Navarro (2007), we can show identification of the components of equations (5) and (6), conditional on a given training history. To apply the theorem,

[14] See also Abbring and Heckman (2007) and Heckman et al. (2016).

we need independent variation (exclusion restrictions). The proof follows a standard identification at infinity argument: use the fact that we can vary the values of X_i^I to find a limit set in which there is no selection bias and thus recover $\beta^I(h_t)$, $\beta^Y(h_t;j)$, and the distribution of unobservables (up to scale) $(\eta_i^I(h_t), \eta_i^Y(h_t;j))$. Note, however, that this result holds conditioning on a particular h_t; as we do not observe individuals in two potential training histories, we cannot identify the joint distribution of outcomes and choices across $h_t \in \mathcal{H}_t$. Therefore, the distribution of joint counterfactuals is not identified without further structure.

To achieve identification of joint counterfactuals, we follow a well-established literature and posit that a low-dimensional set of factors generate dependence across across choices and outcomes for all possible h_t (Aakvik et al., 2005; Carneiro et al., 2003; Heckman et al., 2006, 2016, 2018). Aakvik et al. (2005) and Carneiro et al. (2003) show how researchers can use factor models to identify joint distributions of outcomes by restricting the dependence across unobservables. We assume that unoservables follow:

$$\eta_i^I(h_t) = \lambda^I(h_t)\theta_i + \varepsilon_i^I(h_t), \tag{7}$$

$$\eta_i^Y(h_t;j) = \lambda^Y(h_t;j)\theta_i + \varepsilon_i^Y(h_t;j). \tag{8}$$

where θ_i (with $E(\theta_i) = 0$) represents a fixed, latent ability endowment known by the agent but not the econometrician, $\varepsilon_i^I(h_t)$ is an unobserved measurement error term, and $\epsilon_i^Y(h_t;j)$ is an idiosyncratic shock to productivity that the agent cannot anticipate. Equation (8) imply that a common (and constant) factor, θ_i, drives the endogeneity of choices and outcomes. The rest of the unobserved components of the model are independent across choices and outcomes as well as across time and training histories. This conditional independence assumption extends the matching-on-observables assumption by matching on latent ability. Formally, we assume that $\varepsilon_i^I(h) \perp \varepsilon_i^I(h')$ for all history paths (h, h') with $h \neq h'$, $\varepsilon_i^I(h) \perp \varepsilon_i^Y(h;j)$ for all $(h,j) \in \mathcal{H} \times \{0,1\}$, and $\varepsilon_i^Y(h;j) \perp \varepsilon_i^Y(h';j')$ for all distinctive paths $(h;j)$ and $(h';j')$. Thus, conditional on X_i^Y and X_i^I, θ_i generates all cross-correlations of outcomes and choices across histories.

While we do not directly observe θ_i, we can non-parametrically identify its distribution through a measurement system. Following a large body of literature on static and dynamic treatment effects, we identify the distribution of θ_i using a measurement system observed prior to training participation.[15] Let T_k represent a measure of ability, for $k \in \mathcal{K} \equiv \{1, 2, 3, 4\}$. The measurement system is given by:

$$T_{ik} = X_{ik}^T \beta_k^T + \theta_i \lambda_k^T + \varepsilon_{ik}^T \qquad k \in \mathcal{K}, \tag{9}$$

where X_{ik}^T is a vector of exogenous control variables. The error term ε_{ik}^T is assumed independent of all other error terms in the model. The use of additional information serves for two purposes. First, it is necessary for identification of the model, given the absence of a relatively large number of outcomes per node. Second, by using measures capturing pre-labor market skills, we can interpret unobserved heterogeneity as such and explore heterogeneous treatment effects across this relevant margin.

Given the structure laid out in equations (8) and (9), we can use a factor analysis to identify the joint distribution of potential outcomes $\{Y_i(h_1, j), Y_i(h_2, j)...\}_{h_t \in \mathcal{H}_t}$. The general result is presented in Heckman and Navarro (2007) and Heckman et al. (2016). First, we exploit exclusion restrictions to identify the distribution of errors, conditional on h_t. Second, use (9) and a factor analysis to nonparametrically identify the distribution of θ and thus identify the joint distribution of counterfactuals. In sum, both exclusions restrictions and matching on θ are needed for identification. If we had access to θ we could just use matching on observables and all kinds of treatment effects would be identified. Since we do not, exclusions restrictions are essential in recovering the distributions of unobservables across equations.

To illustrate the result, consider the problem of identifying the dependence of unobservables outcomes and choices in different decision nodes. Concretely, suppose we wish to identify $Cov(\eta_i^Y(h; j), \eta_i^I(h'))$, for $h \neq h'$. Using our structure, and abstracting from observables, we have $Cov(\eta_i^Y(h; j), \eta_i^I(h')) = \lambda^Y(h; j)\lambda^I(h')\sigma_\theta$, where σ_θ is the standard deviation of the distribution of θ, F_θ. To recover all of these terms, consider first the problem of identifying F_θ. This problem follows a standard factor analysis. With access to three measures T_{ik}, we can form $Cov(T_{ik}, T_{ik'}) = \lambda_k^T \lambda_{k'}^T \sigma_\theta$, $Cov(T_{ik'}, T_{ik''}) = \lambda_{k'}^T \lambda_{k''}^T \sigma_\theta$, and $Cov(T_{ik}, T_{ik''}) = \lambda_k^T \lambda_{k''}^T \sigma_\theta$. By normalizing one factor loading $\lambda_k^T = 1$ we can recover $\lambda_{k'}^T, \lambda_{k''}^T$ using ratios of covariances:

[15] This strategy has been applied in various papers on human capital investment decisions. Carneiro et al. (2003), Hansen et al. (2004), Heckman et al. (2006), Heckman et al. (2013), Attanasio et al. (2015), Agostinelli and Wiswall (2016a), and Agostinelli and Wiswall (2016b) are just a few examples.

$$\frac{Cov(T_{ik}, T_{ik'})}{Cov(T_{ik'}, T_{ik''})} = \frac{1}{\lambda^T_{k''}},$$
$$\frac{Cov(T_{ik''}, T_{ik})}{Cov(T_{ik}, T_{ik'})} = \frac{\lambda^T_{k''}}{\lambda^T_{k'}}.$$

Given now that factor loading are identified, we can apply Kotlarski (1967) to identify F_θ and the distribution of measurement errors. Aakvik et al. (2005) and Carneiro et al. (2003) show conditions necessary to identify F_θ. In concrete, we need access to at least three measures if we use only one factor.

Finally, we exploit exclusion restrictions together with the previous result to identify the desired object. Armed with exclusion restrictions, we can identify the marginal distributions of $\eta^Y_i(h; j)$ and $\eta^I_i(h')$. Then we can compute $Var(\eta^Y_i(h; j)) = \lambda^Y(h; j)\sigma^2_\theta$ and $Var(\eta^I_i(h')) = \lambda^I(h')\sigma^2_\theta$ from which we can identify the missing terms $\lambda^Y(h; j)$ and $\lambda^I(h')$. We can proceed in this fashion to identify the general structure of covariances across outcomes, choices, and decision nodes.

Once we add observables, we need to assume that latent ability is orthogonal to all observed variables in vector $\mathbf{X}_i \equiv (X^Y_i, X^I_i, \{X^T_{ik}\}_{k\in\mathcal{K}})$. While this assumption is strong, it allows us to capture dimensions of skills that are orthogonal to those who are already partly captured in vector $\mathbf{X}_i$.

As discussed, identification is achieved through independent variation (exclusion restrictions) in the choice equations. In our context, we use the average training hours at the firm and all firms within a certain geographical location ("comuna") where the individual is currently working as node-specific instruments. An implicit assumption behind these exclusion restrictions is that an individual does not alter her behavior—in a way that could affect her earnings—due to working in a firm that is more or less likely to invest in training.[16] We argue that the exclusion restriction may hold in our setting due to various reasons. First, since among FT-participants in 2002-2010, the average worker took up fewer than 20 hours of FT-subsidized training, these courses are unlikely to represent a major consideration for firm-switching decisions among workers. This argument is reinforced by the fact that average training hours at the firm/comuna are not publicly available

[16] Related to this issue, Ba et al. (2017) discuss how identification of the effects of training programs in an experimental setting breaks down when individuals anticipate having a subsidized training in future periods and so they change behavior today (for example, by lowering their present employment search intensity).

information. The exclusion restriction may also be violated if firms with unobserved characteristics which lead them to both invest in off-the-job training tend to hire more productive workers. While we cannot directly test for this hypothesis, we note that the results presented in Table 3 show that including firm fixed-effects does not alter the estimated returns to training. As a result, it does not appear that workers sort into firms due to the firm's policy on FT-course provision.

4.2 Estimation

For estimation purposes, we define the sample likelihood as follows. Let Ψ be the vector that collects the set of parameters. Given our independence assumptions, the likelihood for a set of $\mathcal{I}$ individuals is given by:

$$\mathcal{L}(\Psi \mid \cdot) = \prod_{i \in \mathcal{I}} \left[\int_{\theta} \prod_{k \in \mathcal{K}} f_T(T_{ik} \mid X_{ik}^T, \theta) \prod_{t \in \mathcal{T}} \varphi(Y_i \mid X_i^Y, X_i^I, h_t, \theta) dF(\theta) \right],$$

where $f_{T_k}(\cdot)$ is the conditional density function of test score k, $F(\theta)$ represents the unobserved ability's cumulative distribution function, and

$$\begin{aligned} \varphi(Y_i \mid X_i^Y, X_i^I, h_t, \theta) \quad = \quad & \left[f(Y_i(h_t; 1) \mid X_i^Y, \theta) \Pr(I_i(h_t) \geq 0 \mid X_i^I, \theta) \right]^{D_i(h_t)} \\ & \times \left[f(Y_i(h_t; 0) \mid X_i^Y, \theta) \Pr(I_i(h_t) < 0 \mid X_i^I, \theta) \right]^{1 - D_i(h_t)}, \end{aligned}$$

with $f_Y(\cdot)$ representing the conditional density function of Y_i. We use normal distributions for the idiosyncratic shocks in the choice process (equation 5), earnings regression (equation 6), and measurement system (equation 9). Even though we assume normal disturbances, note that our identification argument does not rely on normality and we only assume it for computational convenience. However, for estimation purposes, we do adopt a flexible functional form for $F(\theta)$:

$$\theta \sim \rho_1 N(\tau_1, \sigma_1^2) + \rho_2 N(\tau_2, \sigma_2^2).$$

We estimate the model by Markov Chain Monte Carlo (MCMC). Using the estimated model, we simulate 20 samples from the original sample, each new sample associated to a different draw from the posterior of distribution of structural parameters.

Table 4 shows the variables we include in the measurement system, training probit and earnings equation. In the measurement system, we include age at the time at the time of PSU, as it may affect test score performance. To identify the distribution of the latent factor, we rely on four different measures, including language and math PSU test scores, high school GPA, and initial wage at the time of entry. While the literature on latent factors has traditionally relied on test score measures to identify unobserved ability, these papers largely consider human capital investment in the context of schooling choices. Our context is different, as we examine post-schooling training choices. To this end, we include the initial wage in the measurement system, departing from the existing literature, to better capture baseline productivity prior to training choices. In the choice equations, we use age and, as noted above, we include the average training hours at the firm and all firms within a certain geographical location (*comuna*) where the individual is currently working. Lastly, in the earnings equation, we include a gender dummy, age, and a constant.

Tables B.1-B.3 show the estimated parameters of our econometric model. In Table B.1, we present the estimated parameters of the measurement system. We find that latent ability loads positively on both test score measures, high school GPA, and the initial salary. To understand the relative contribution of observed characteristics and latent ability vector to test score measures, we present a variance decomposition in Figure 2. The unobserved ability measure explains 64% of the variance in the initial salary, but only 10%, 6% and 4% of the math PSU and verbal PSU scores and high school GPA, respectively. Hence, our measure of unobserved ability is more related to initial labor market productivity than workers' academic performance at high school graduation. We note that including the initial salary as part of the measurement system involves a trade-off in estimating the distribution of ability: it reduces the share of the variance in test scores explained by the unobserved factor but allows us to better capture unobserved ability at the time of labor market entry. In Figure 3, we depict the distribution of the unobserved ability factor. As the estimated unconditional distribution of θ does not exhibit considerable deviations from normality, we note that assuming a standard normal distribution for the distribution of unobserved ability should not significantly change selection patters and the estimated returns to training.

Tables B.2 and B.3 present the estimated parameters of the training and earnings equations. Across all choice nodes, women are more likely to participate in training and more skilled individuals are more likely to participate in job training. Moreover, workers in firms with a large share of workers participating in FT courses as well as those in geographic areas with more training course availability are more likely to have participated in training in any period.[17] The earnings equations indicate that males outearn women by upwards of 0.07 log points. No discernible pattern arises with respect to the age-earnings profile. Latent ability has positive effects on earnings, and this result holds across all training nodes.

Lastly, we assess the model's accuracy in predicting observed outcomes and choices. Figure 4 compares observed and simulated training histories. The model matches training decisions well, both in the first and second year. Table 5 contrasts the means and standard deviations of log wages by year and training choices. Overall, simulated earnings show some differences with observed earnings but these gaps are smaller than 0.07 log points in all but one case.

4.3 Selection on Unobserved Characteristics

Figure 5 compares the density of the unobserved ability for workers choosing different training paths. Training choices are denoted by $(h_2; j)$, where h_2 represents the training history prior to $t = 2$ (or the first period decision) and j captures the training choice at $t = 2$. We find significant differences in the density of ability across one-time participants, depending on the timing of the decision. While the latent ability for those who choose training in the first period but not in the second $(1; 0)$ almost entirely overlaps with the density of the never-trained group, the distribution of the latent factor for those only trained in the second period clearly surpasses the never-trained group. The always-trained group $(1; 1)$ has the largest ability stock relative to that of workers following other training paths, resembling differences shown in baseline test scores in Table 2. In fact, the never-trained group trails their always-trained counterparts by 0.83 standard deviations in the latent skill distribution. Overall, we find evidence of sorting on unobservables, as higher-skilled workers are more likely to have participated in job training.

5 Returns to Job Training

This section presents evidence on the impact of job training on earnings. We define and estimate static and dynamic treatment effects. We examine the mechanisms driving the dynamic effects of

[17] The only statistically insignificant coefficient corresponds to the one on local-level course availability in the training equation of period $t = 2$, conditional on training in the first period.

training, by estimating dynamic complementarity (substitutability) in the context of job training participation as well as continuation values. We also examine heterogeneous impacts across the latent ability distribution.

5.1 Static Treatment Effects

We first present evidence on static treatment effects, which capture the effects of training conditional on reaching a particular choice node. Given that we examine earnings in the quarter following the training event, this parameter recovers the short-term effects of training. Let $E[.]$ denote the expected value taken with respect to the distribution of $(\mathbf{X}, \theta, \varepsilon)$, where ε is the collection of idiosyncratic shocks determining outcomes and choices $\varepsilon \equiv (\varepsilon_j^I, \varepsilon_s^Y)$. We first present evidence on the average treatment effect in period t, ATE_t, defined as the average impact of period t training on period t earnings, conditional on a training history h_t. Formally,

$$ATE(h_t) \equiv E[Y_i(h_t; 1) - Y_i(h_t; 0)], \tag{10}$$

We can additionally defined the average effect of training in period t conditional on having participated in t given h_t. That is, the treatment on the treated, $TT(h_t)$, parameter is defined as:

$$TT(h_t) \equiv E[Y_i(h_t; 1) - Y_i(h_t; 0) \mid D_i(h_t) = 1]. \tag{11}$$

We present the estimated static returns to job training in Table 6. The average short-term returns to first-period training ($ATE(h_1)$) equal 1.7%, which are largely similar to the corresponding treatment on the treated parameter for individuals who in fact took up first-year training. The estimated returns are far lower than those found in the OLS regression presented in Table 3, since the reduced form approach fails to account for sorting on unobserved characteristics. For second-period training, conditional on not training in the first year ($h_2 = \{0\}$), we find larger static returns, reaching 3.4%. This effect exceeds that for second-year participation for workers who had been trained in the first period ($h_2 = \{1\}$), which equals 0.4%. Across both second-period returns, we find that the treatment on the treated parameters are in the same order of magnitude with the average treatment effect parameters. Lastly, we note that while the ATE parameters are positive on average, a considerable share of workers would enjoy a negative return from program participation. For instance, first-period training lowers earnings for 45% of workers in our sample.[18]

The estimated short-run effects of training (ATEs and TTs) are similar to the fixed-effects estimates, presented in Table 3. However, these parameters do not necessarily coincide with the parameters identified by fixed-effects regressions. The coefficient associated with job training in the fixed-effects model is a weighted average—with potentially negative weights— across the three treatment effect parameters (Callaway and Sant'Anna, 2018; de Chaisemartin and D'Haultfœuille, 2018; Goodman-Bacon, 2018). Furthermore, as we show in Web Appendix A, the conditions for a first-difference estimator to recover ATE are not met in our sample.

Figures 6 and 7 examine heterogeneous effects of training across the ability distribution. To compute these figures, we simulate outcomes and choices drawing different values from the distribution of theta.[19] For each simulation we compute the individual-level treatment effect and then estimate a non-linear regression of the estimated treatment effects onto the latent skill distribution as a way to summarize this relationship. In the first year of training (Figure 6), we find larger returns to participation for less-skilled workers, though the returns profile is largely flat across the skill distribution. Larger differences appear in the returns to second year training (Figure 7). Among workers who had not been initially trained ($h_2 = \{0\}$), training has larger effects on earnings for high-skill workers, exceeding 4% for those in the top decile of the skill distribution, while remaining close to 3% for individuals in the bottom decile. For first-period participants ($h_2 = \{1\}$), the impact of second-period training for workers above the median of the latent skill distribution is not different from zero, while exceeding 1% for those in the top ability decile.[20]

In short, the estimated static treatment effects reveal heterogeneous impacts across different

[18] We also find negative earnings returns for a non-negligible share of program participants. Since our econometric model is agnostic about the role of expectations in the decision-making process, we cannot directly distinguish whether the negative returns could be explained through financial regret—individuals do not correctly predict the monetary gains following from training—or through psychic costs—the agent is willing to accept a negative monetary return because training yields non pecuniary benefits.

[19] Linearity in the relationship between individual-level ATE and TT and ability is not mechanically imposed by our linear-in-the-parameters outcome equations. As we force people from not training to training, and we move across the ability distribution, choices in the second period might change, which could introduce non-linearities in the θ-ATE relationship.

[20] In two training nodes, job training has larger effects on less-skilled individuals. This type of reversed-Roy selection has also been found in other contexts. Kline and Walters (2016) find that children who are less likely to attend Head Start gain more from it, which is consistent with evidence from Germany (Cornelissen et al., 2017). Mountjoy (2018) finds suggestive evidence that the returns to two-year colleges are larger for those who are less likely to attend, a result in line with findings in the context of post-secondary education in Chile (Rodriguez et al., 2016).

decision margins, job training histories, and latent ability. As such, the constant effect framework required in fixed-effect estimators is rejected in favor of a model of differential short-term gains from job training.

5.2 Dynamic Treatment Effects

While we have so far analyzed short-term treatment effects, job training may also affect medium- and long-term labor market outcomes. In this sub-section, we extend our analysis by estimating the dynamic returns to training, continuation values and dynamic complementarity/substitutability. As such, we adapt the framework introduced in Section 2 to the two period setting. Thus,

$$\widetilde{Y}_i(h_1; j) \equiv Y_i(h_1; j) + \rho\left(D_i(h_2)Y_i(h_2; 1) + (1 - D_i(h_2))Y_i(h_2; 0)\right), \qquad j \in \mathcal{H}_2 \equiv \{0, 1\}$$

where $D_i(h_2)$ denotes second-period participation and $Y_i(h_2; j)$ captures earnings for training choice j, given history h_2. The long-term direct effect introduced in equation (2) represents the direct effect of training on earnings two years after the event. Meanwhile, the continuation value of training recovers the additional gain (if any) of training in the second-period from training in first period for those who take up training in $t = 2$.

In Table 7, we present estimates from the following dynamic treatment effect parameters:

$$DATE \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0)\right], \tag{12}$$

$$DTT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 1\right], \tag{13}$$

$$DTUT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 0\right], \tag{14}$$

and decompose them into short- and medium-term direct effects and continuation values.[21] The dynamic average treatment effect ($DATE$) indicates that job training in the first period in the labor force increases the present value of earnings by 3.6%. This result is driven largely by the direct effect of job training, as first-period participation increases earnings (in present-value terms) by 1.1 and 2.5% one and two years after training, respectively. However, early training lowers the returns to training in the second period—the continuation value—by 0.04%, though the effect is

[21]For this exercise, we assume a discount factor $\rho = 1/(1.05)$. In addition, we present our results as % increase from the average baseline present value of earnings.

not statistically significant. This result is consistent with the estimates of Table 6, which indicate a lower return to second-period training, for those who were trained in the first period compared to their non-trained counterparts. All in all, the direct effects make up for the bulk of the estimated $DATE$ (101%), whereas continuation value accounts for a small part of the total effect. As with the static treatment effects, we find similar effects in the dynamic TT and TUT parameters.

In Figure 8, we examine heterogeneous dynamic treatment effects across the latent skill distribution, and decompose them into direct effects and the continuation value. Confirming the results shown in Figure 6, the short-term direct effects of training are decreasing across the skill distribution. In contrast, the medium-term direct effects are larger for more skilled workers. Lastly, we find that the continuation value of training decreases across the ability distribution. The combination of these three elements leads to a $DATE$ parameter that is largely flat across the ability distribution.

We note that the estimated continuation value of job training stands in contrast with those estimated in the context of formal schooling. For example, Heckman et al. (2016) find that the bulk of the return to high school graduation (around 70%) and to college enrollment (25%) is explained by the continuation value of schooling, reaching a larger share for more skilled students. On the other hand, we find negative continuation values from job training for high-ability workers, yet the contribution of continuation values to the total return to first-period training is small compared to that of the direct effect components.

Dynamic Complementarity (Substitutability). The relationship between continuation values and dynamic complementarity (substitutability) in the two-period model can be expressed as follows:

$$\begin{aligned}
&\underbrace{E\left[D_i(1)(Y_i(1;1)-Y_i(1,0))-D_i(0)(Y_i(0;1)-Y_{i2}(0;0))\right]}_{\text{Continuation Value}} = \\
&\underbrace{E\left(Y_i(1;1)-Y_i(1;0)\right)-\left(Y_i(0;1)-Y_i(0;0)\right)]}_{\text{Dynamic Complementarity/Substitutability}} + \\
&\underbrace{E\left(D_i(1)-1\right)(Y_i(1;1)-Y_i(1;0))-(D_i(0)-1)(Y_i(0;1)-Y_i(0;0))]}_{\text{Dynamic Sorting Gains}} \qquad (15)
\end{aligned}$$

As a result, on average, the production function of job training exhibits dynamic complementarity (substitutability) if the return from training in a second period is higher (lower) conditional on first-period participation: $E\left[Y_i(1;1)-Y_i(1;0)\right] > E\left[Y_i(0;1)-Y_i(0;0)\right]$.

In Figure 9, we present evidence from a local polynomial regression of dynamic complementarity (substitutability) parameter onto the unobserved ability distribution (equation (15)). First-period job training lowers the return from subsequent participation, independent of second-period decisions: on average, we find evidence of dynamic substitutability (-2.7%), which increases with ability. On the other hand, dynamic sorting gains are positive (2.6% on average) and increasing with ability (3% for those in the top ability decile), which fits in with the sorting patterns documented in Section 4.3.

We note that dynamic substitutability, in the context of job training, may arise for various reasons. First, job training could comprise multiple courses covering topics in the same area, with workers starting in a baseline course and subsequently taking part in more complex coursework.[22] In this setting, dynamic substitutability may appear if the first course delivers critical information for improving job performance, with subsequent courses delivering less value-added. As a result, while early trainees would take the second course in their second year in the labor force, non-trainees would participate in the initial course, which delivers larger returns—thus yielding dynamic substitutability. This result may also appear in a context of course heterogeneity, with individuals choosing the most important (or higher-return) courses early on in their labor market careers and subsequently taking courses delivering lower returns.[23] Workers could rationally follow such a strategy as the returns to the high-payoff courses could be enjoyed over a longer time horizon (Ben-Porath, 1967).

Dynamic substitutability could also emerge if workers face the decision to either accumulate human capital within the firm, or "outside" the firm, through job training courses. Since the returns to training would then recover the gains from formal off-the-job training relative to on-the-job training, dynamic substitutability would indicate increasing returns to within-firm learning over time, rather than capturing the underlying technology of job training. In this context, the returns to training should be higher for workers switching firms, as the human capital accumulated in job training would transfer to the new employer, though this would not be the case for prior within-firm learning. To test for this possibility, we replicate our empirical analysis for a subsample of workers

[22] For example, workers first need to learn to operate a computer before taking a course on a specific software.

[23] In this context, workers would first take courses directly related to their industry or occupation and subsequently participate in foreign language courses, for example.

who do not switch firms in their first two years in the labor force ("stayers") in Web Appendix C.[24] As we find similar static and dynamic returns to training in the stayer subsample (Tables C.2 and C.3), we argue that firm-switching behavior induced through job training participation does not drive the estimated impacts of job training. The results presented in this section indicate small, yet significant gains arising from job training. We next examine whether these gains would be actionable upon for workers, by considering policy relevant treatment effects.

6 Dynamic Policy-Relevant Treatment Effects

We have so far focused on estimating parameters such as the average treatment effect and the treatment on the treated. However, these may not necessarily be relevant parameters for policy purposes. For instance, workers induced to change their training choices through a particular policy may have different observed and unobserved characteristics relative to the average worker, and thus their estimated gains from training would not be captured by average population parameters (Mogstad and Torgovitsky, 2018). In this section, we introduce a framework which allows us to examine the returns to various policy alternatives, decompose the effects across response types, and study dynamic responses to these policies.

6.1 Policy Intervention

In this context, we follow the literature which defines policy relevant treatment effects in terms of policy shocks that do not affect marginal treatment effects (Heckman and Vytlacil, 2001; Carneiro et al., 2010; Mogstad et al., 2018; Mogstad and Torgovitsky, 2018), by analyzing the effect of a policy that affects the net cost of first-period training while leaving fixed the net cost of second-period training. Concretely, we examine the effect of an increase in the number of average hours of FT courses available across *comunas* in $t = 1$ only. In practice, our simulation may capture a temporary, unexpected shock to the training industry that increases the number of available courses in the market through a policy intervention.[25] Even though the policy change only affects

[24]This sample consists of 22,247 of the original 37,089 workers—that is, 60% of our sample never switched firms in the first two years of labor force participation. We do not directly model workers' firm switching behavior, yet remark that the characteristics and training choices of stayers are largely similar to those in the full sample (Figure C.1 and Table C.1).

[25]This policy change could also take place through a subsidy for FT providers to develop additional courses for first-year labor market participants. Our framework can be extended to the case of policies that shift the cost of

net costs of job training in the first period directly, it alters second-period decisions by shifting workers' progression through the training tree depicted in Figure 1. Furthermore, the simulated policy change affects observed outcomes exclusively through training choices, not by influencing counterfactual earnings.

To consider how such a policy would impact training choices, we introduce the following notation. Let $D_i^a(h_t)$ be the training choice in period t in a given state of the world a. We model the policy change as a shift from a to a', which may directly result in changed training decisions in the first period. For instance, workers who are first-period compliers are characterized by $\{D_i^{a'}(h_1) = 1,\ D_i^a(h_1) = 0\}$.[26] The policy change may also affect second-period choices through changes in first-period decisions, as first-period compliers have reached a different choice node. Counterfactual outcomes are otherwise unaltered: $Y_i^a(h_t; j) = Y_i^{a'}(h_t; j) = Y_i(h_t; j)$, $j \in \{0, 1\}$.[27]

Both the number of workers affected by the policy and the estimated earnings effects might depend on the magnitude of the intervention, which calls into question the external validity of the Local Average Treatment Effect ($LATE$) of a particular policy change. Mogstad and Torgovitsky (2018) discuss the importance of considering the external validity of estimated treatment effect parameters, and they suggest extrapolating local average treatment effects as a solution for the limited external validity of $LATE$s. To this end, we incorporate this consideration in our policy simulation by estimating the effect of a policy intervention of varying sizes, simulating a 10 and 50 percent expansion in the number of FT-hours available in each *comuna* in the first period.[28]

6.2 Counterfactual Choices and Outcomes

Observed earnings in both periods depend on training choices. Let $Y_{i,t}^a$ be observed earnings in period t (after training decision was made) under a. Given our assumption about the nature of the

training across multiple time periods. Here, we study the simplest case to illustrate the benefits of our model in terms of revealing dynamic policy responses.

[26] The simulated policy implies that a represents the current state of the world, whereas a' captures increased training availability in the first-period and baseline second-period course availability.

[27] In each policy state a, we keep fixed draws of all error terms and parameters from the baseline model. Thus, differences in choices between a and a' stem exclusively from changes in the net utility of training participation through a shift in the local availability of course hours. The simulated policy affects the latent utility associated with the first participation node. As such, whether the policy impacts' workers participation decisions depends on the coefficient associated with the instrument being shifted in the policy simulation.

[28] Since the baseline number of hours per worker in each comuna in the first year equals 0.55 hours, the 10% increase equals an increase in 0.055 course hours, while the 50% increase results in an average increase of 0.275 hours.

policy, in $t = 1$:

$$Y_{i,1}^a \equiv D_i^a(h_1)Y_i(h_1;1) + (1 - D_i^a(h_1))Y_i(h_1;0). \tag{16}$$

where, again, we note that the policy change only affects first-period earnings for workers changing their initial training decision. A similar expression, but encompassing the sequence of decisions defined by $D_i^a(h_1)$ and $D_i^a(h_2)$, can be obtained for $Y_{i,2}^a$.

We are interested in the effects of the policy on the present value of earnings, so let $\widetilde{Y}_i^a \equiv Y_{i,1}^a + \rho Y_{i,2}^a$, where ρ is the discount factor. The effect of a policy that shifts the net benefits of training choices from a to a' is given by $E\left[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^a\right]$. Since the policy only affects this parameter through changes in training choices rather than through a direct impact on earnings, we can decompose its effect by identifying agents' response types. The policy intervention may lead workers to change their first-period participation decision (compliers and defiers) or to maintain their baseline choice under policy state a (always-takers and never-takers). Given our implicit monotonicity assumption in equation (5), an increase in first-period course availability will only affect outcomes through first-period compliers—captured by $D_i^a(h_1) = 0$ and $D_i^{a'}(h_1) = 1$. However, this change would may also affect second-period choices, as workers who modified their first-period decision due to the policy change may make different training choices depending on their training history. Therefore, the group of first-period compliers can be further divided by second-period responses in four types: compliers-always takers, compliers-compliers, compliers-never takers, and compliers-defiers.[29,30] As the policy change does not have an impact on counterfactual earnings, the effect for workers not changing their training decision in either period will not be different from zero. All in all, the effect of the policy change on the present value of earnings is given by $\Delta_{a,a'} = E\left[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^a \mid D_i^{a'}(h_1) = 1, D_i^a(h_1) = 0\right]$, the dynamic policy relevant treatment parameter ($DPRTE$). If we let $\mathbb{A}_i$ denote the condition $(D_i^{a'}(h_1) = 1, D_i^a(h_1) = 0)$, this parameter can be

[29] We note than an alternative policy change affecting the utility of second-period training participation could also impact the policy parameter of interest through workers who did not change their initial participation decision, but who became second-period compliers. As such, estimating the effects of this policy intervention requires considering the impacts on initial period always-takers and never-takers, who became second-period compliers. Of course, evaluating policies in a dynamic context contrasts policy-relevant analysis in an static world, where only contemporary compliers are the relevant group (Heckman and Vytlacil, 2001; Mogstad and Torgovitsky, 2018).

[30] In a similar set-up, Heckman et al. (2016) decompose $LATE$ into the effects of augmenting the availability of colleges on earnings for different subgroups affected by the policy in a dynamic-discrete choice model. While their analysis considers earnings impacts in one time period for groups who shift their previous choices, we instead analyze how both choices and outcomes of different periods are affected by the policy.

decomposed as:

$$
\begin{aligned}
\Delta_{a,a'} = & \underbrace{E[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^{a} | D_i^{a'}(1) = D_i^{a}(0) = 1, \mathbb{A}_i]}_{\text{Compliers, Always-Takers } (W_{CO,AT})} \times \underbrace{\Pr[D_i^{a'}(1) = D_i^{a}(0) = 1 | \mathbb{A}_i]}_{\text{Probability of Always-Takers}} \\
& + \underbrace{E[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^{a} | D_i^{a'}(1) = 1, D_i^{a}(0) = 0, \mathbb{A}_i]}_{\text{Compliers, Compliers } (W_{CO,CO})} \times \underbrace{\Pr[D_i^{a'}(1) = 1, D_i^{a}(0) = 0 | \mathbb{A}_i]}_{\text{Probability of Compliers}} \\
& + \underbrace{E[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^{a} | D_i^{a'}(1) = 0, D_i^{a}(0) = 0, \mathbb{A}_i]}_{\text{Compliers, Never-Takers } (W_{CO,NT})} \times \underbrace{\Pr[D_i^{a'}(1) = 0, D_i^{a}(0) = 0 | \mathbb{A}_i]}_{\text{Probability of Never-Takers}} \\
& + \underbrace{E[\widetilde{Y}_i^{a'} - \widetilde{Y}_i^{a} | D_i^{a'}(1) = 0, D_i^{a}(0) = 1, \mathbb{A}_i]}_{\text{Compliers, Defiers } (W_{CO,DF})} \times \underbrace{\Pr[D_i^{a'}(1) = 0, D_i^{a}(0) = 1 | \mathbb{A}_i]}_{\text{Probability of Defiers}} \qquad (17)
\end{aligned}
$$

In this set-up, we can therefore estimate the aggregate effect of the policy on the net present value of earnings for affected workers and examine the impacts across dynamic response groups. For example, the increase in FT-course availability induces compliers-compliers to move from never taking up job training to participating in the two periods. The complier-defier group may arise if a particular sub-set of workers induced to participate in the first period would take up training in the second period had they not been early trainees. These groups might reveal policy-relevant behavior and we directly test for their presence in our empirical analysis. Note that the weights are given by the prevalence of each response type as a share of all workers who change participation decisions due to the policy.

Note that, in our dynamic model, extrapolation following the existing MTE literature cannot be pursued. Suppose we have as many instruments for different potential choices given by the number of potential training histories individuals can take. In our simple empirical example of two choices in two moments in time, one can estimate a 2SLS model onto three different dummies (given a baseline choice) which are instrumented with three different instruments (the original two plus their interaction). In this multinomial model, IV will not be able to separately identify effects of training for different compliers types, that is, the components of equation (6.2) (Heckman and Urzúa, 2010; Kirkeboen et al., 2016; Mountjoy, 2018). Thus the need for our structural model to identify different dynamic groups for which hypothetical policies have heterogeneous effects.

6.3 Results

The first panel of Table 8 presents the share of workers induced to change their training decisions due to the policy change. We find that a 10% increase in course-hour availability would induce 0.5% of the workers in our sample to participate in job training at some point in their first two years in the labor force. The share of "affected" workers expands linearly across program expansion size, since a 50 percent increase in FT course-hours would induce 2.4% of young workers to change their training decision.[31] Moreover, we find that almost half of all policy compliers come from the complier-never taker group, which captures workers who take up job training only in the first period in response to the policy. Meanwhile, the complier-complier group, who take-up training in both time periods due to the policy change, accounts for 30% of all affected workers. The weights assigned to dynamic response types are largely constant across the two interventions.

In the second panel of Table 8, we show the estimated effect of the policy simulation on the present value of earnings. We find that a small increase in course availability would increase the earnings of affected workers by 3.1%, reaching 3.5% for a 50 percent increase in course-hour availability. While the impacts are captured by the $DPRTE$ parameter, the effect sizes fit largely in line with the dynamic average treatment effect ($DATE$) presented in Table 7.

We find that the simulated policy would have heterogeneous impacts across dynamic response types. For instance, workers who are induced to take-up training in both periods (compliers-compliers) enjoy an earnings increase in the 4% range. We find similar effects for workers who are induced to participate in *only* one training course, reaching 3.2% for compliers-never-takers. The complier-defier group, which captures workers who move up their training choice due to the policy change, faces heterogeneous earnings gains by program size expansion. For small increases in course-hour availability, their earnings would increase by 1.8%, though this effect is not statistically significant. On the other hand, earnings would increase by 2.7% under the larger simulated expansion.

In Figure 10, we explore whether differences in the density of the unobserved factor across response type groups may account for the heterogeneous earnings impacts documented above. We find that the latent ability density of compliers-always-takers dominates that of the other response

[31] While the program expansion need not have linear effects on training take-up, the empirical evidence shows no indication of non-linear responses by program expansion size.

types, surpassing the ability of compliers-compliers by 0.11 standard deviations, on average, providing evidence that individuals respond to the policy based on the latent ability. Moreover, there are larger differences with the other response types, as the average latent ability of compliers-always takers exceeds that of compliers-never takers by 0.42 standard deviations. In this context, we examine whether the 50 percent expansion in first-year course hours would deliver heterogeneous returns by estimating a local polynomial regression of the dynamic policy relevant treatment effect parameter against the distribution of latent ability (Figure 11).[32] We find largely homogeneous effects across the ability distribution, in the range of 3%.

In sum, we have found that increased course availability would lead a non-negligible share of workers to change their early-career training decisions, while positively affecting their medium-term earnings. However, the heterogeneous impacts across dynamic response types suggest that the positive impacts depend on the number and the timing of the courses that workers are induced to take up. Since the static framework necessarily overlooks multi-period response types, these results highlight the importance of considering dynamic effects of policy changes.

7 Conclusions

In this paper, we leverage a large government-subsidized program to present the first estimates of repeated participation in job training for first-time labor market entrants. We document dynamic selection patterns and show that high ability workers are more likely to participate in job training early in their careers, though among one-time participants, second-period trainees are more skilled vis-a-vis those trained in the first year. We find that the static returns to job training are positive and significant, though they vary across the timing of the event, training histories and heterogeneously across the latent ability distribution. The dynamic treatment effects indicate larger medium-term gains from early job training, and this effect is fully explained through the direct effect of training. As such, the continuation value of training is not different from zero, standing in contrast with the positive continuation value from schooling (Heckman et al., 2016). We further document the differences between continuation values and dynamic complementarity

[32] The optimal policy design depends on the cost function of course hour expansion. While it would be natural to think that costs follows a convex pattern, the Chilean government has recently allowed e-learning courses to be included as part of *Franquicia Tributaria* courses, suggesting that a larger program expansion need not have a larger per unit cost than a small expansion (SENCE, 2015).

(substitutability). We find dynamic substitutability of first-year training: early investments decrease the economic returns to later investments. The result is stronger for high-skilled workers. Dynamic substitutability may be explained by the structure of the job training courses examined in this paper, or more generally by the structure of post-schooling human capital accumulation processes. While we cannot test the potential mechanisms formally, we consider our results a first step towards understanding the complex dynamic of the returns to training.

Moreover, while estimating a variety of treatment effects allows us to capture the various margins through which training affects labor market outcomes, these returns may not necessarily be actionable upon for workers (Mogstad and Torgovitsky, 2018). We have therefore examined the estimated impacts of an expansion in course-hour availability for first-time labor market entrants. In this context, we identify dynamic response types and estimated dynamic policy relevant treatment effects, as early-career policy changes may affect workers' subsequent training decisions. While the increase in course availability would lead to a sizable increase in medium-term earnings, the effects are heterogeneous across dynamic response types. As a result, we remark that any policymaker considering an expansion in training courses should take into account the potential impact on workers' subsequent labor market trajectories, rather than focusing solely on short-term outcomes.

References

Aakvik, A., J. J. Heckman, and E. J. Vytlacil (2005, mar). Estimating treatment effects for discrete outcomes when responses to treatment vary: an application to Norwegian vocational rehabilitation programs. *Journal of Econometrics 125*(1-2), 15–51.

Abbring, J. and G. van den Berg (2003). The non-parametric identification of treatment effects in duration models. *Econometrica 45*, 1491–1517.

Abbring, J. H. and J. J. Heckman (2007). Econometric evaluation of social programs, part iii: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. *Handbook of econometrics 6*, 5145–5303.

Acemoglu, D. and D. Autor (2011). Chapter 12 Skills, Tasks and Technologies: Implications for Employment and Earnings. In D. Card and O. Ashenfelter (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 1043–1171.

Agostinelli, F. and M. Wiswall (2016a). Estimating the Technology of Children's Skill Formation. NBER Working Paper No. 22442.

Agostinelli, F. and M. Wiswall (2016b). Identification of Dynamic Latent Factor Models: The Implication of Re-Normalization in a Model of Child Development. NBER Working Paper No. 22441.

Albert, C., C. Garcia-Serrano, and V. Hernanz (2010). On-the-job training in europe: Determinants and wage returns. *International Labour Review 149*(3), 315–341.

Andersson, F., H. Holzer, J. Lane, D. Rosenblum, and J. Smith (2013). Does federally-funded job training work? Nonexperimental estimates of WIA training impacts using longitudinal data on workers and firms. NBER Working Paper No. 19446.

Arellano, M. and B. Honore (2001). Panel Data Models: Some Recent Developments. In J. Heckman and E. Lazear (Eds.), *Handbook of Econometrics*.

Ashenfelter, O. (1978). Estimating the effect of training programs on earnings. *The Review of Economics and Statistics 60*(1), 47–57.

Ashenfelter, O. and D. Card (1985). Using the longitudinal structure of earnings to estimate the effect of training programs. *The Review of Economics and Statistics 67*(4), 648–660.

Attanasio, O., S. Cattan, E. Fitzsimons, C. Meghir, and M. Rubio-Codina (2015). Estimating the Production Function For Human Capital: Results From a Randomized Control Trail in Colombia. NBER Working Paper No. 20965.

Autor, D. H., F. Levy, and R. J. Murnane (2003, nov). *The Quarterly Journal of Economics 118*(4), 1279–1333.

Ba, B., J. C. Ham, R. J. Lalonde, and X. Li (2017). Estimating (easily interpreted) dynamic training effects from experimental data. *Journal of Labor Economics 35*(S1).

Barnow, B. S. and J. Smith (2016). Employment and Training Programs. In R. A. Moffitt (Ed.), *Economics of Means-Tested Transfer Programs in the United States, Volume II*. Chicago: The University of Chicago Press.

Ben-Porath, Y. (1967). The Production of Human Capital and the Life Cycle of Earnings. *Journal of Political Economy 75*, 352–365.

Booth, A. L. (1993). Private sector training and graduate earnings. *The Review of Economics and Statistics 75*, 164–170.

Callaway, B. and P. H. C. Sant'Anna (2018). Difference-in-Differences With Multiple Time Periods and an Application on the Minimum Wage and Employment. DETU Working Papers 1804, Department of Economics, Temple University.

Card, D., J. Kluve, and A. Weber (2010, nov). Active Labour Market Policy Evaluations: A Meta-Analysis. *The Economic Journal 120*(548), F452–F477.

Carneiro, P., K. T. Hansen, and J. J. Heckman (2003, may). 2001 Lawrence R. Klein Lecture Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College Choice. *International Economic Review 44*(2), 361–422.

Carneiro, P., J. J. Heckman, and E. Vytlacil (2010, jan). Evaluating Marginal Policy Changes and the Average Effect of Treatment for Individuals at the Margin. *Econometrica 78*(1), 377–394.

Carrillo, F., M. E. De Diego, and R. Miranda (2018). Caracterización del Funcionamiento de la Franquicia Tributaria. Unpublished Manuscript, Comisin Nacional de Productividad.

Cattaneo, M. D. (2010). Efficient semiparametric estimation of multi-valued treatment effects under ignorability. *Journal of Econometrics 155*(2), 138–154.

Comisión Revisora del Sistema de Capacitación e Intermediación Laboral (2011). Informe Final.

Cornelissen, T., C. Dustmann, A. Raute, and U. Schönberg (2017). Who benefits from universal childcare? Estimating marginal returns to early childcare attendance. Working Paper.

Cunha, F., J. J. Heckman, L. Lochner, and D. V. Masterov (2006). Interpreting the Evidence on Life Cycle Skill Formation. In E. Hanushek, S. Machin, and L. Woessmann (Eds.), *Handbook of the Economics of Education*, Volume 1 of *Handbook of the Economics of Education*, pp. 697–812. Elsevier.

de Chaisemartin, C. and X. D'Haultfœuille (2018). Two-way fixed effects estimators with heterogeneous treatment effects. Working Paper.

Dyke, A., C. J. Heinrich, P. R. Mueser, K. R. Troske, and K.-S. Jeon (2006). The effects of welfare-to-work program activities on labor market outcomes. *Journal of Labor Economics 24*(3), 567–607.

Fitzenberger, B., M. Furdas, and C. Sajons (2016). End-of-year spending and the long-run employment effects of training programs for the unemployed. *IZA DP No. 10441*.

Fitzenberger, B. and R. Völter (2007, aug). Long-run effects of training programs for the unemployed in East Germany. *Labour Economics 14*(4), 730–755.

Frazis, H. and M. A. Loewenstein (2007). *On-the-Job Training*. Now Publishers, Inc.

Fredriksson, P. and P. Johansson (2008). Dynamic treatment assignment: The consequences for evaluations using observational data. *Journal of Business and Economic Statistics* (4).

Gelbach, J. B. (2016). When do covariates matter? and which ones, and how much? *Journal of Labor Economics 34*(2), 509–543.

Goodman-Bacon, A. (2018). Difference-in-Differences with Variation in Treatment Timing. Working Paper.

Guvenen, F., B. Kuruscu, S. Tanaka, and D. Wiczer (2015). Multidimensional skill mismatch. *Federal Reserve Bank of Minneapolis. Research Department Working Paper 729*.

Hansen, K. T., J. J. Heckman, and K. J. Mullen (2004). The effect of schooling and ability on achievement test scores. *Journal of Econometrics 121*(1), 39–98.

Heckman, J., H. Ichimura, J. Smith, and P. Todd (1998). Characterizing Selection Bias Using Experimental Data. *Econometrica 66*(5), 1017.

Heckman, J., R. Pinto, and P. Savelyev (2013, oct). Understanding the Mechanisms Through Which an Influential Early Childhood Program Boosted Adult Outcomes. *American Economic Review 103*(6), 2052–2086.

Heckman, J. J. (1981). Statistical models for discrete panel data. In C. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Panel Data with Econometric Applications*, pp. 114–78. MIT Press.

Heckman, J. J. and V. J. Hotz (1989). Choosing among alternative nonexperimental methods for estimating the impact of social programs: The case of manpower training. *Journal of the American statistical Association 84*(408), 862–874.

Heckman, J. J., J. E. Humphries, and G. Veramendi (2016). Dynamic treatment effects. *Journal of Econometrics 191*(2), 276–292.

Heckman, J. J., J. E. Humphries, and G. Veramendi (2018). Returns to education: The causal effects of education on earnings, health, and smoking. *Journal of Political Economy 126*(S1), S197–S246.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *The review of economic studies 64*(4), 605–654.

Heckman, J. J., R. J. LaLonde, and J. A. Smith (1999). The Economics and Econometrics of Active Labor Market Programs. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labour Economics: Volume 4B*, pp. 1865–2097. Elsevier Masson SAS.

Heckman, J. J. and S. Navarro (2007). Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics 136*(2), 341–396.

Heckman, J. J., J. Stixrud, and S. Urzua (2006, jul). The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior. *Journal of Labor Economics 24*(3), 411–482.

Heckman, J. J. and S. Urzúa (2010). Comparing IV with structural models: What simple IV can and cannot identify. *Journal of Econometrics 156*(1), 27–37.

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

Heckman, J. J. and E. Vytlacil (2005). Structural equations, treatment effects, and econometric policy evaluation 1. *Econometrica 73*(3), 669–738.

Ingram, B. F. and G. R. Neumann (2006). The returns to skill. *Labour Economics 13*(1), 35–59.

Johnson, R. C. and C. K. Jackson (2017, June). Reducing inequality through dynamic complementarity: Evidence from head start and public school spending. Working Paper 23489, National Bureau of Economic Research.

Keane, M. P., P. E. Todd, and K. I. Wolpin (2011). Chapter 4 The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics*, Volume 4, pp. 331–461.

Kirkeboen, L. J., E. Leuven, and M. Mogstad (2016, aug). Field of Study, Earnings, and Self-Selection. *The Quarterly Journal of Economics 131*(3), 1057–1111.

Kline, P. and C. R. Walters (2016, nov). Evaluating Public Programs with Close Substitutes: The Case of Head Start. *The Quarterly Journal of Economics 131*(4), 1795–1848.

Kotlarski, I. I. (1967). On Characterizing the Gamma and the Normal Distribution. *Pacific Journal of Mathematics 20*(1).

Larsson, L. (2003). Evaluation of swedish youth labor market programs. *Journal of human resources 38*(4), 891–927.

Lechner, M. (2000). An evaluation of public-sector-sponsored continuous vocational training programs in east germany. *Journal of Human Resources*, 347–375.

Lechner, M. and C. Wunsch (2009). Are training programs more effective when unemployment is high? *Journal of Labor Economics 27*(4), 653–692.

Lengermann, P. (1999). How long do the benefits of training last? evidence of long term effects across current and previous employers. *Research in Labor Economics 19*.

Lise, J., C. Meguir, and J.-M. Robin (2016)). Mismatch, sorting, and wages. *Review of Economic Dynamics 53*(19), 63–87.

Lynch, L. M. (1992). Private-sector training and the earnings of young workers. *American Economic Review 82*, 299–312.

Mincer, J. (1974). Schooling, experience, and earnings. human behavior & social institutions no. 2.

Mogstad, M., A. Santos, and A. Torgovitsky (2018). Using Instrumental Variables for Inference about Policy Relevant Treatment Effects. Forthcoming in *Econometrica*.

Mogstad, M. and A. Torgovitsky (2018). Identification and Extrapolation with Instrumental Variables. Forthcoming in *Annual Review of Economics*.

Mountjoy, J. (2018). Community Colleges and Upward Mobility. Unpublished manuscript.

Mueser, P. R., K. R. Troske, and A. Gorislavsky (2007). Using state administrative data to measure program performance. *The Review of Economics and Statistics 89*(4), 761–783.

Postel-Vinay, F. and J. Lise (2015). Multidimensional skills, sorting, and human capital accumulation. *Meeting Papers 386*.

Rodriguez, J., S. Urzua, and L. Reyes (2016, apr). Heterogeneous Economic Returns to Post-Secondary Degrees: Evidence from Chile. *Journal of Human Resources 51*(2), 416–460.

Saltiel, F., M. Sarzosa, and S. Urzua (2018). Cognitive and Socio-emotional Ability. In G. Jones, J. Johnes, T. Agasisti, and L. Lopez-Torres (Eds.), *Handbook of Contemporary Education Economics*. Edwar Elgar Publishing.

Sanders, C. and C. Taber (2012, sep). Life-Cycle Wage Growth and Heterogeneous Human Capital. *Annual Review of Economics 4*(1), 399–425.

SENCE (2015). Resolucion exenta numero 1664.

Smith, J. A. and P. E. Todd (2005). Does matching overcome lalonde's critique of nonexperimental estimators? *Journal of econometrics 125*(1-2), 305–353.

SpitzOener, A. (2006, apr). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics 24*(2), 235–270.

Veum, J. (1997). Training and job mobility among young workers in the united states. *Journal of Population Economics 10*, 219–233.

Table 1: Summary Statistics

Variable	Mean	(Std. Dev.)
Male	0.46	(0.5)
Age at Graduation	17.81	(0.64)
Math PSU (Standardized)	-0.04	(0.98)
Verbal PSU (Standardized)	-0.03	(0.99)
High School GPA (Standardized)	0.01	(0.99)
Monthly Salary after First Year (USD)	551.04	(333.4)
Monthly Salary after Second Year (USD)	623.96	(388.31)
Observations	37,089	

Notes: Table 1 presents summary statistics of our estimation sample (see Section 3). The dependent variable is the monthly average of earnings in the first quarter following each training stint. For simplicity, we refer to this variable as concurrent with the training decision. Tests scores (Math and Verbal) and high school GPA are standardized across the general population of test-takers to be of mean zero and variance 1.

Table 2: Summary Statistics by Training Node

	h_3 = (0,0)	h_3 = (0,1)	h_3 = (1,0)	h_3 = (1,1)
Male	0.46	0.46	0.47	0.43
Age at Graduation	17.82	17.79	17.81	17.78
Math PSU (Standardized)	-0.10	0.04	-0.04	0.19
Verbal PSU (Standardized)	-0.07	0.02	-0.03	0.17
High School GPA (Standardized)	-0.01	0.04	0.00	0.14
Monthly Salary after First Year (USD)	506.00	611.15	566.32	732.64
Monthly Salary after Second Year (USD)	572.32	702.36	640.29	820.12
Observations	23675	5306	4360	3748

Notes: Table 2 presents summary statistics of the estimation sample (see Section 3) across different training histories. Training histories after two periods are given by $h_3 = (h, h')$, where h, h' denote training choices for period $t = 1$ and $t = 2$ ($h, h' \in \{0, 1\}$), respectively.

Table 3: Reduced-Form Estimates: Returns to Job Training

Panel A. Short-Term Returns to Job Training

	First-Period Earnings (Y_{i1})		Second Period Earnings (Y_{i2})			
	(1)	(2)	(3)	(4)	(5)	(6)
D_{i1}	0.182***	0.149***				
	(0.007)	(0.006)				
$D_{i2}(1)$			0.257***	0.161***		
			(0.013)	(0.011)		
$D_{i2}(0)$					0.209***	0.178***
					(0.008)	(0.007)
Control Variables		X		X		X
Observations	37089	37089	8108	8108	28981	28981

Panel B. Returns to Job Training: Panel Data Estimates

	(1)	(2)	(3)	(4)
Effect of training	0.223***	0.179***	0.009***	0.013***
	(0.005)	(0.004)	(0.003)	(0.002)
OLS	X			
OLS + controls		X		
Individual FE			X	
Individual and firm FE				X
Observations ($N \times T$)	74178	74178	74178	66352

Notes: Table 3 presents regressions of log-earnings against a dummy variable capturing job training participation. Control variables include college entrance exam performance, high school GPA and age. The dependent variable is the monthly average of earnings in the first quarter following the training period. Panel A presents short-term returns to training, examining how training in period t affects Y_{it}. Panel B considers the effects of training on earnings, exploiting the longitudinal component of the data. Column (1) presents OLS regressions without control variables. Column (2) includes PSU test scores, highschool GPA, a gender dummy, age, and age squared. Column (3) computes the first differences estimator. Column (4) uses an individual and firm fixed-effect. p-values are in parenthesis, where * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

Table 4: Variables Used in Implementation of the Model

Variables	Earnings Equation	Training Probit	Measurement System
Constant	Yes	Yes	Yes
Gender	Yes	Yes	Yes
Age at Test			Yes
Age in Year t	Yes	Yes	
Cognitive Factor	Yes	Yes	Yes
Average Training Hours at Firm		Yes	
Average *Comuna* Hours		Yes	

Notes: Table 4 shows the variables used in our empirical model. In the measurement system, we use math and language college entrance test scores, high school GPA and the initial salary and include as the observed measures. Training decisions include gender and age as control variables as well as training-course availability across training decision nodes.

Table 5: Goodness of Fit: Labor Market Outcomes by Training History

Estimate	$Y_i(h_2 = \{1\})$	$Y_i(h_2 = \{0\})$	$Y_i(h_3 = \{1,1\})$	$Y_i(h_3 = \{1,0\})$	$Y_i(h_3 = \{0,1\})$	$Y_i(h_3 = \{0,0\})$
A. Means						
Actual	6.30	6.12	6.55	6.29	6.40	6.20
	[6.29,6.32]	[6.12,6.13]	[6.53,6.57]	[6.28,6.31]	[6.39,6.42]	[6.19,6.20]
Model	6.23	6.14	6.42	6.27	6.35	6.23
B. Standard Deviations						
Actual	0.57	0.52	0.58	0.57	0.55	0.54
Model	0.51	0.50	0.53	0.52	0.53	0.52

Notes: Table 5 shows the means of log earnings by year and training choice from the observed data and the simulated sample. $Y_i(h_2 = \{h\})$ denotes earnings after period $t = 1$, where h represents the training decision in the first period. $Y_i(h_3 = \{j, j'\})$ represents earnings following $t = 2$, with $h, h' \in 0, 1$ recovering first and second period training decisions, respectively. In brackets, we show a 95% confidence interval on the mean of observed earnings. The second panel shows the observed and estimated standard deviation of earnings across training histories.

Table 6: Static Returns to Job Training (in %)

Treatment effect	$t = 1$	$t = 2$	
		$h_2 = \{0\}$	$h_2 = \{1\}$
ATE (percentage points)	1.69	3.39	0.43
	[1.65,1.72]	[3.31,3.47]	[0.27,0.58]
TT (percentage points)	1.79	3.51	0.35
	[1.72,1.86]	[3.31,3.70]	[0.11,0.60]
$Pr(ATE < 0) \times 100$	45.33	45.89	49.46
	[45.21,45.44]	[45.76,46.02]	[49.22,49.71]
$Pr(TT < 0) \times 100$	45.14	45.66	49.55
	[44.90,45.39]	[45.37,45.96]	[49.17,49.93]

Notes: Table 6 presented the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters along with the share of workers who enjoy negative returns to training across training histories and over time. The first column shows the returns to first-period training ($t = 1$). The second and third columns present the estimated returns to second-period job training, conditional on not participating in first-period training (Column 2) and conditional on first-period participation (Column 3). We show 95% confidence intervals in brackets.

Table 7: Dynamic Returns to First-Period Job Training (in %)

	DATE	DTT	DTUT
Direct Effect (short-term)	1.11	1.16	1.10
	(0.01)	(0.02)	(0.01)
	[31%]	[32%]	[31%]
Direct Effect (medium-term)	2.51	2.49	2.52
	(0.02)	(0.05)	(0.02)
	[70%]	[69%]	[70%]
Continuation Value	-0.04	-0.05	-0.04
	(0.02)	(0.04)	(0.02)
	[-1%]	[-1%]	[-1%]
Total	3.58	3.59	3.58
	(0.02)	(0.05)	(0.03)

Notes: Table 7 presents the estimated Dynamic Average Treatment Effects (DATE), Dynamic Treatment on the Treated (DTT) and Dynamic Treatment on the Untreated (DTUT) of first-period training on the present value of earnings. Let $\widetilde{Y}_i(h_1; j) \equiv Y_i(h_1; j) + \rho\, (D_i(h_2)Y_i(h_2; 1) + (1 - D_i(h_2))Y_i(h_2; 0))\; j \in \mathcal{H}_2 \equiv \{0, 1\}$ be the present value of earnings for training choice $j \in \{0, 1\}$. The individual-level dynamic treatment effect equals $\widetilde{Y}_i(1) - \widetilde{Y}_i(0)$, where

$$\widetilde{Y}_i(1) - \widetilde{Y}_i(0) = \underbrace{(Y_i(1) - Y_i(0))}_{\text{Direct effect (short-term)}} + \underbrace{\rho[Y_i(0; 1) - Y_i(0; 0)]}_{\text{Direct effect (medium-term)}} + \underbrace{\rho\left[D_{i2}(1)(Y_i(1; 1) - Y_i(0; 1)) - D_{i2}(0)(Y_i(1; 0) - Y_i(0; 0))\right]}_{\text{Continuation value}}.$$

We show the above decomposition for:

$$DATE \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0)\right],$$
$$DTT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 1\right],$$
$$DTUT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 0\right],$$

We present DATE, DTT, and DTUT as percentage of mean baseline of the present value of earnings ($E[\widetilde{Y}_i(0)]$). We present standard errors in parenthesis and the percentage contribution of each term in brackets.

Table 8: Share of Compliers and Dynamic-Policy Treatment Effects (in %)

Panel A. Share of Compliers by Intervention Size and Dynamic Response Type Group Weights

Share Compliers		
	Policy: +10% (1)	Policy: +50% (2)
Compliers (CO)	0.47%	2.38%
	Weights by type	
Compliers-Always Takers (CO, AT)	0.106	0.105
Compliers-Compliers (CO, CO)	0.297	0.294
Compliers-Never Takers (CO, NT)	0.489	0.487
Compliers-Defiers (CO, DFT)	0.106	0.109

Panel B. Dynamic Policy Relevant Treatment Effect by Response Type

Dynamic Policy Relevant Treatment Effects		
	Policy: +10% (1)	Policy: +50% (2)
Compliers	3.09%	3.51%
	[2.44,3.74]	[3.23,3.79]
Compliers-Always Takers (CO, AT)	1.54%	2.05%
	[-0.39,3.46]	[1.20,2.91]
Compliers-Compliers (CO, CO)	3.97%	4.02%
	[2.81,5.14]	[3.51,4.53]
Compliers-Never Takers (CO, NT)	3.19%	3.71%
	[2.24,4.14]	[3.30,4.12]
Compliers-Defiers (CO, DF)	1.84%	2.65%
	[-0.18,3.87]	[1.78,3.51]

Notes: Table 8 shows policy-relevant treatment effects of two policy shocks: a temporary increase in FT-hours by 10 and 50%. We present the effect of the policy on the present value of earnings conditional on being a W_j type of complier (as a percentage of baseline earnings), a 95% confidence interval of these returns (in brackets), and the proportion of compliers (in % terms). Formally, let $\widetilde{Y}_i^{a_j} \equiv Y_{i1}^{a_j} + \rho Y_{i2}^{a_j}$ be the present value of earnings for a policy regime a_j. Following the notation used in the main text, this table presents the probabilities associated with the different groups (weights) and associated earnings changes, where

$$CO, AT \equiv \left\{ D_i^{a'}(h_1) = 1, D_i^{a}(h_1) = 0, D_i^{a'}(1) = 1, D_i^{a}(0) = 1 \right\},$$

$$CO, CO \equiv \left\{ D_i^{a'}(h_1) = 1, D_i^{a}(h_1) = 0, D_i^{a'}(1) = 1, D_i^{a}(0) = 0 \right\},$$

$$CO, NT \equiv \left\{ D_i^{a'}(h_1) = 1, D_i^{a}(h_1) = 0, D_i^{a'}(1) = 0, D_i^{a}(0) = 0 \right\},$$

$$CO, DF \equiv \left\{ D_i^{a'}(h_1) = 1, D_i^{a}(h_1) = 0, D_i^{a'}(1) = 0, D_i^{a}(0) = 1 \right\}.$$

$$CO \equiv \left\{ D_i^{a'}(h_1) = 1, D_i^{a}(h_1) = 0 \right\}.$$

Figure 1: Observed Training Choices

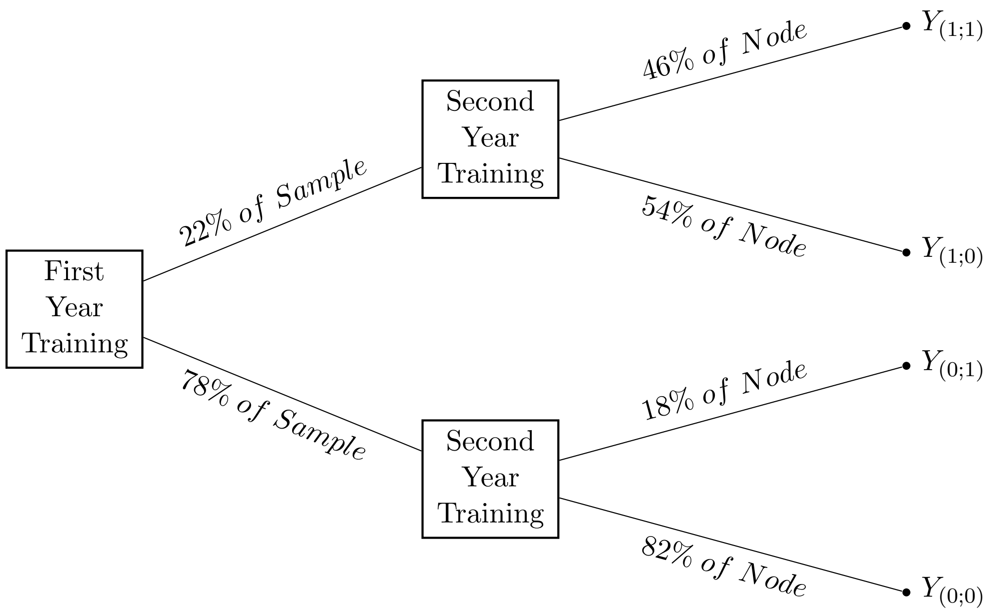

Note: In Figure 1, we present the decision tree through which workers decide whether to participate in training in each of their first two years in the labor force. In each node, we include the observed share of workers in our sample who decide to participate in training.

Figure 2: Measurement System: Variance Decomposition

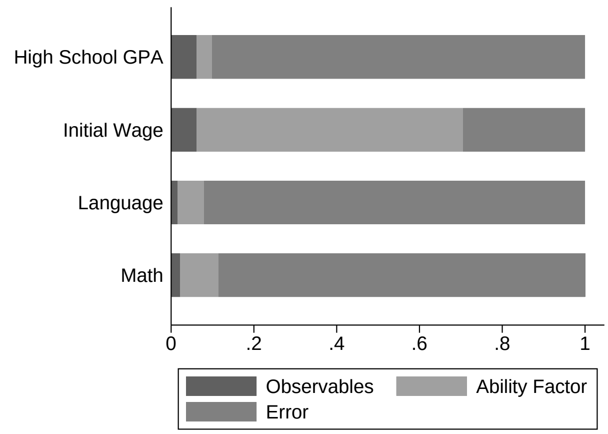

Note: In Figure 2, we show the contribution of each variable to the variance of observed measures using the simulated sample from our model. The "Observables" row indicates the share of the variance of the measure explained by the observed variables: age at the time of test score, gender, parental employment dummy variables, mother's and father's education, as well as household size. The "Ability Factor" component shows the proportion of the test score variance explained by unobserved ability. Finally, the "Error term" represents the share of the variance in each observed measure explained by the unobserved idiosyncratic error of the measurement equation.

Figure 3: Distribution of Unobserved Ability

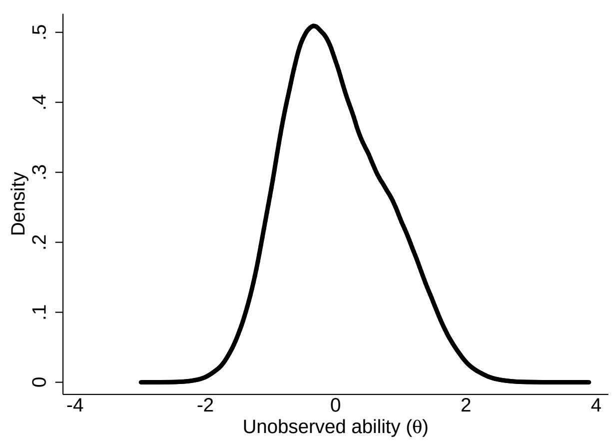

Note: In Figure 3, we show the estimated density of the unobserved ability factor. We obtain this density using the simulated sample from our estimated model. We approximate the distribution of the individual's unobserved ability factor by a mixture of normal distributions.

Figure 4: Goodness of Fit: Training Decisions

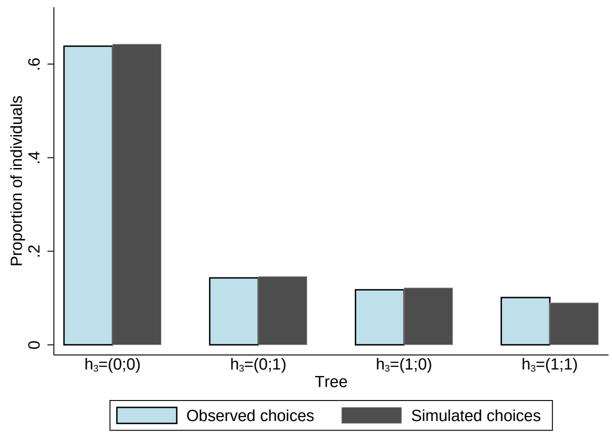

Note: In Figure 4, we compare the share of workers who followed each of the four possible training histories in their first two years in the labor force. A training history is given by $h_3 = (h; h')$, where the first and second entry indicate training decisions in the first and second period, respectively.

Figure 5: Distribution of Unobserved Ability by Training History

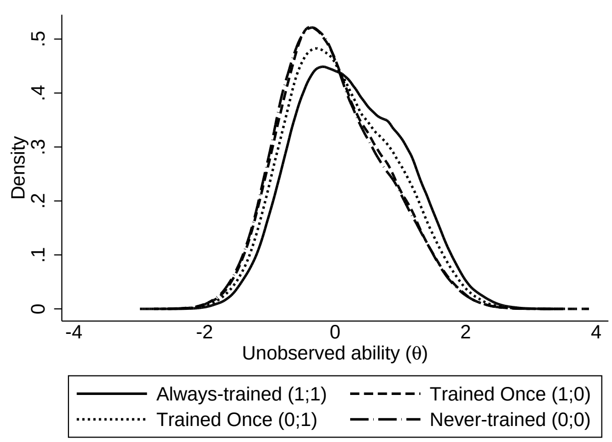

Note: Figure 5 shows the estimated density of unobserved ability for different training paths for workers in their first two years in the labor force. A training history is defined by $h_3 = (h; h')$, where h and h' denotes training choices $(h, h' \in \{0, 1\})$ for periods 1 and 2, respectively.

Figure 6: Heterogeneous Returns to First-Period Training Participation

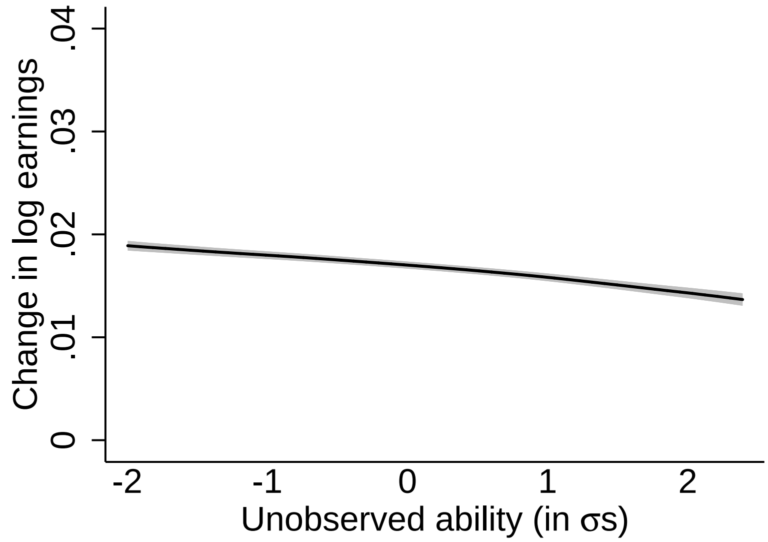

(a) Average Treatment Effect

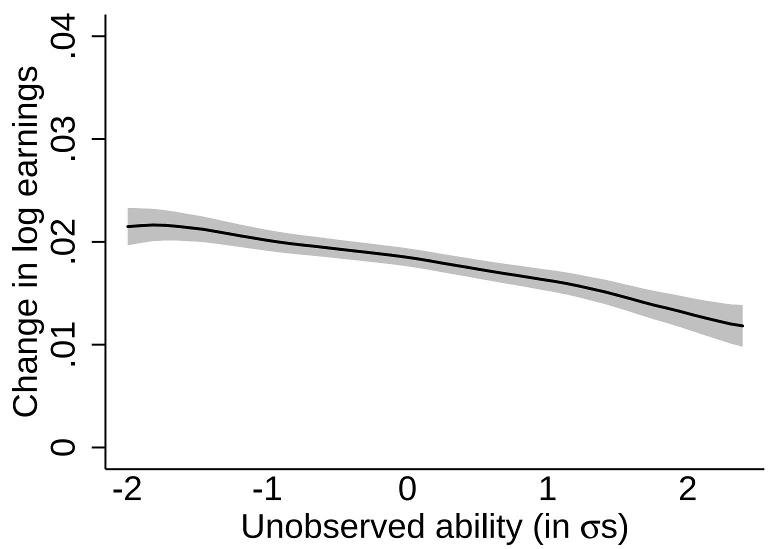

(b) Treatment on the Treated

Note: In Figure 6, we estimate local polynomial regressions of the estimated ATE and TT parameters for the first training event (at $t = 1$) against the distribution of unobserved ability.

Figure 7: Heterogeneous Returns to Second-Period Training Participation

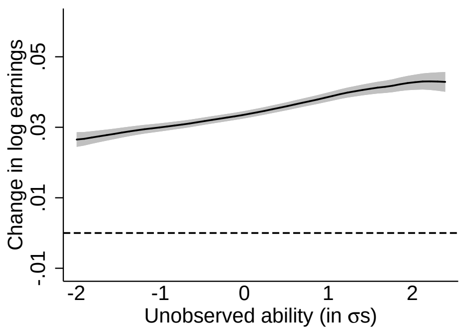

(a) ATE ($h_2 = 0$)

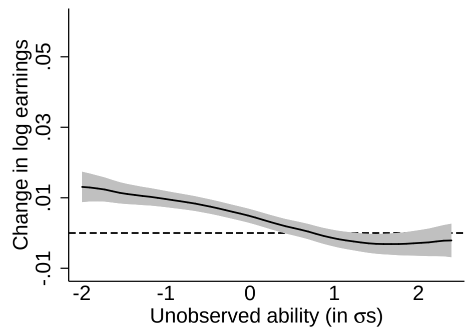

(b) ATE ($h_2 = 1$)

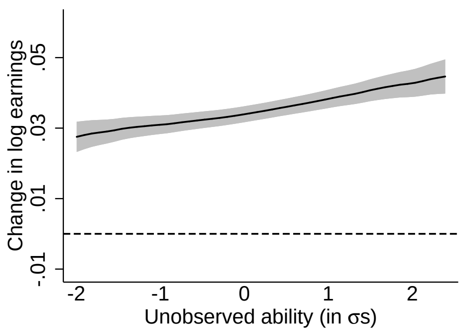

(c) TT ($h_2 = 0$)

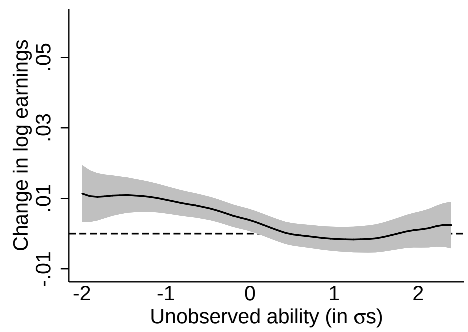

(d) TT ($h_2 = 1$)

Note: In Figure 7, we estimate local polynomial regressions of the estimated average treatment effect and treatment on the treated parameters of second-period job training participation against the distribution of latent ability, conditional on first-period choices ($h_2 = j$ where $j \in \{0, 1\}$). The upper-left panel shows, for instance, the average treatment effect of second-period participation for workers who had not taken up training in the first year.

Figure 8: Heterogeneous Dynamic Returns to First-Period Training Participation

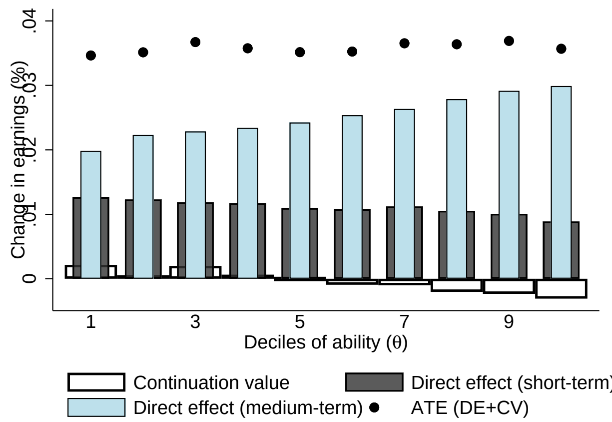

(a) DATE

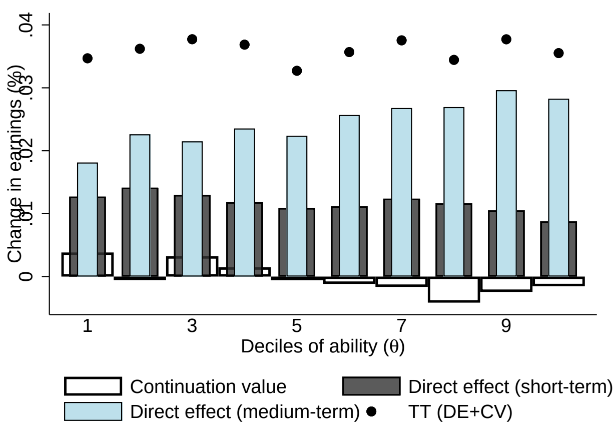

(b) DTT

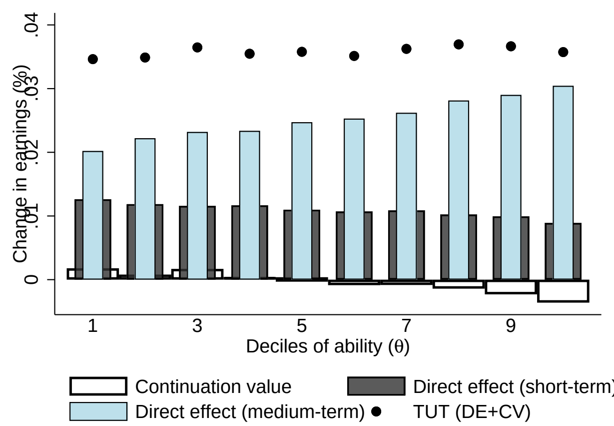

(c) DTUT

Note: In Figure 8, we estimate Dynamic Average Treatment Effects (DATE), Dynamic Treatment on the Treated (DTT) and Dynamic Treatment on the Untreated (DTUT) of training in $t = 1$ on the present value of earnings across deciles of unobserved ability. Let $\widetilde{Y}_i(h_1; j) \equiv Y_i(h_1; j) + \rho\,(D_i(h_2)Y_i(h_2; 1) + (1 - D_i(h_2))Y_i(h_2; 0))\, j \in \mathcal{H}_2 \equiv \{0, 1\}$ be the present value of earnings for training choice $j \in \{0, 1\}$. The individual-level dynamic treatment effect equals $\widetilde{Y}_i(1) - \widetilde{Y}_i(0)$, where

$$\widetilde{Y}_i(1) - \widetilde{Y}_i(0) = \underbrace{(Y_i(1) - Y_i(0))}_{\text{Direct effect (short-term)}} + \underbrace{\rho[Y_i(0;1) - Y_i(0;0)]}_{\text{Direct effect (medium-term)}} + \underbrace{\rho\,[D_{i2}(1)(Y_i(1;1) - Y_i(0;1)) - D_{i2}(0)(Y_i(1;0) - Y_i(0;0))]}_{\text{Continuation value}}.$$

We show the above decomposition for:

$$DATE \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0)\right],$$

$$DTT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 1\right],$$

$$DTUT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 0\right],$$

We present DATE, DTT, and DTUT as percentage of mean baseline of the present value of earnings ($E[\widetilde{Y}_i(0)]$). We present standard errors in parenthesis and the percentage contribution of each term in brackets.

Figure 9: Heterogeneous Dynamic Complementarity (Substitutability) and Dynamic Sorting

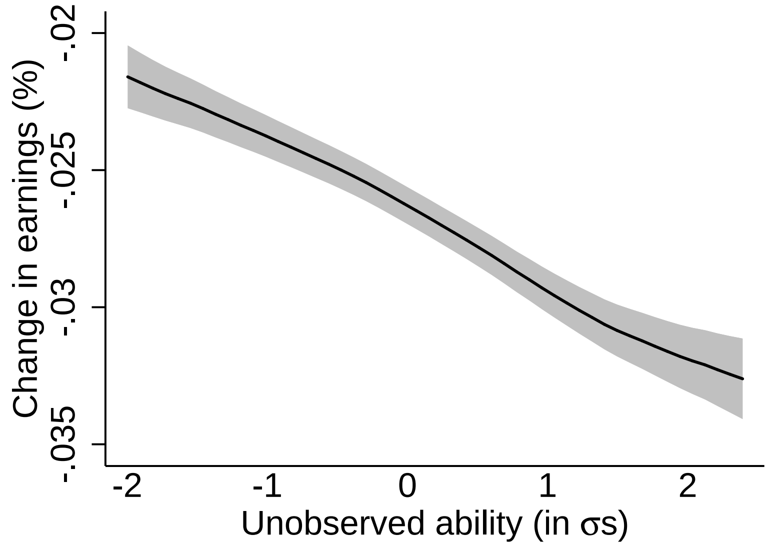

(a) Dynamic Complementarity (Substitutability)

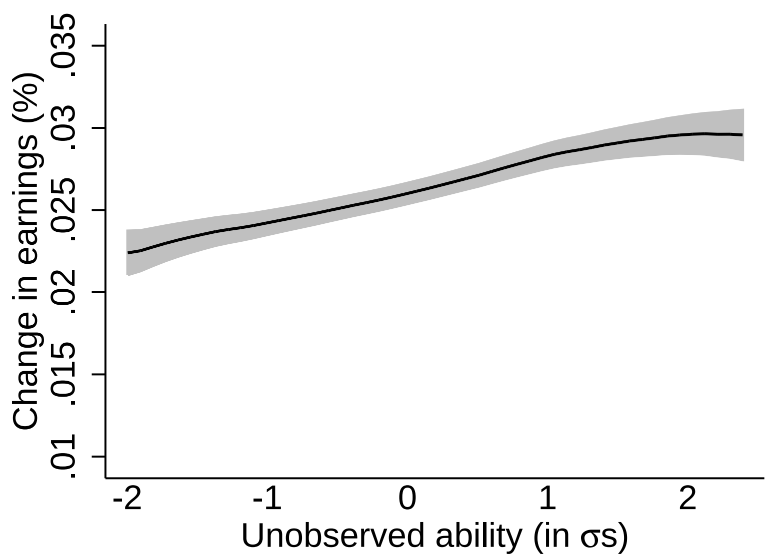

(b) Dynamic Sorting Gains

Note: In Figure 9, we show dynamic complementarity (substitutability) and dynamic sorting gains (as a percentage of mean baseline of the present value of earnings) as a function of unobserved ability, where

$$\underbrace{D_i(1)(Y_i(1;1) - Y_i(1,0)) - D_i(0)(Y_i(0;1) - Y_{i2}(0;0))}_{\text{Continuation Value}} = \underbrace{(Y_i(1;1) - Y_i(1;0)) - (Y_i(0;1) - Y_i(0;0))}_{\text{Dynamic Complementarity/Substitutability}} + \underbrace{(D_i(1) - 1)(Y_i(1;1) - Y_i(1;0)) - (D_i(0) - 1)(Y_i(0;1) - Y_i(0;0))}_{\text{Dynamic Sorting Gains}}$$

Figure 10: Density of Unobserved Ability by Dynamic Response Types

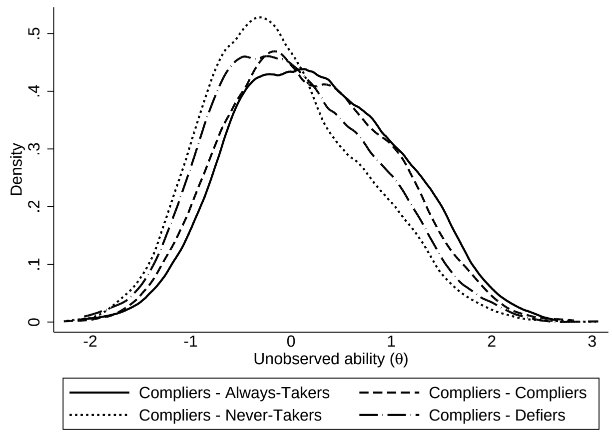

Note: Figure 10 shows the estimated density of unobserved ability across dynamic response types. The simulated policy reflects a 50% expansion in program course-hours availability in both periods. The density of the latent factor across response types is similar across different program expansion levels.

Figure 11: Heterogeneous Dynamic Policy Relevant Treatment Effects

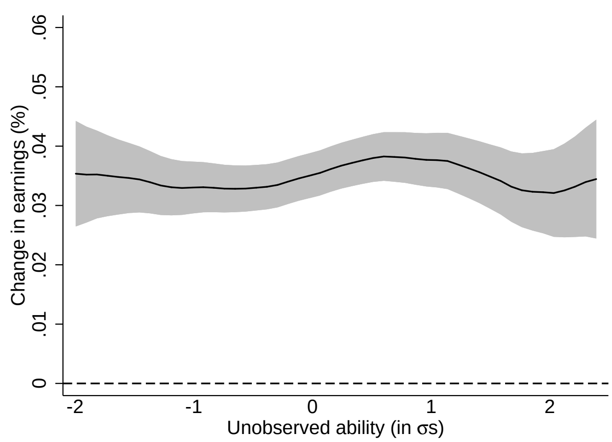

Note: Figure 11 shows the estimated impact of a 50% expansion in course availability on the present-value of earnings across levels of unobserved ability, captured by the dynamic policy-relevant treatment effect parameter.

Appendices

A FD Estimator and Treatment Effects

Here, we analyze if fixed-effects estimators recover the ATE of training choices. We show that, in the context of our model, the fixed-effect estimator can solve the inconsistency problem of OLS by controlling for unobserved heterogeneity. However, it fails to identify the average treatment effect.

By using the structure of our model, we formally define the ATE. The impact of training for an individual i at period t for a given history h equals $Y_{it}(h_t;1) - Y_{it}(h_t;0)$.[33] Let $H_{it}(h_t)$ be an indicator variable that equals 1 if individual i in period t followed training history h and 0 otherwise. The overall average of these individual treatment effects is defined as:

$$\begin{aligned} ATE &\equiv E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h_t)\left(Y_{it}(h_t;1) - Y_{it}(h_t;0)\right)\right] \\ &= E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h_t)\left(\mu^Y(h_t;1) - \mu^Y(h_t;0) + (\lambda^Y(h_t;1) - \lambda^Y(h_t;0))\theta_i\right)\right], \end{aligned} \tag{A.1}$$

where the expected value operator integrates with respect to i and t. Therefore, ATE is a weighted average of individual treatment effects across periods and different potential training histories.

In a longitudinal data set-up, the analyst's goal is to identify (A.1) using observed data (Y_{it}, D_{it}), where D_{it} and Y_{it} represent the observed training indicator and outcome variable. As a starting point, consider the following linear regression:

$$Y_{it} = \pi_0 + \pi_1 D_{it} + \xi_{it} \qquad \text{for } i = 1, ..., N \text{ and } t = 1, ..., T \tag{A.2}$$

where ξ_{it} is an error term. OLS identifies:

$$\delta^{OLS} \equiv \frac{Cov(Y_{it}, D_{it})}{Var(D_{it})} = E[Y_{it}|D_{it} = 1] - E[Y_{it}|D_{it} = 0]$$

If the data generating process follows our dynamic model, then potential self-selection into training results in a correlation between ξ_{it} and D_{it} (Ashenfelter and Card, 1985). To see how self-selection affects the reduced-form estimate, first, let us define the following:

$$\mu^Y(j) \equiv \sum_{h\in\mathcal{H}_t} H_{it}(h_t)\mu^Y(h_t,j), \quad \lambda^Y(j) \equiv \sum_{h_t\in\mathcal{H}_t} H_{it}(h_t)\lambda^Y(h_t,j), \quad \epsilon^Y_{it}(j) \equiv \sum_{h_t\in\mathcal{H}_t} H_{it}(h_t)\epsilon^Y_{it}(h_t,j)$$

for $j \in \{0,1\}$. Second, following the standard switching regression model, we can express observed variables (Y_{it}, D_{it}) as functions of underlying potential outcomes and choices. Observed variables

[33] We note that the notation in this section differs slightly from the main text, as we include the time period in which earnings are observed, defined as Y_{it}.

are given by:

$$D_{it} \equiv \sum_{h_t \in \mathcal{H}_t} H_{it}(h_t) D_{it}(h_t), \tag{A.3}$$

$$Y_{it} \equiv \sum_{h_t \in \mathcal{H}_t} H_{it} \left[D_{it}(h_t) Y_{it}(h_t; 1) + (1 - D_{it}(h_t)) Y_{it}(h_t; 0) \right]. \tag{A.4}$$

Using the above definitions of D_{it} and Y_{it}, and summing over observed and unobserved parameters across training histories, we have:

$$Y_{it} = \mu^Y(0) + D_{it}(\mu^Y(1) - \mu^Y(0)) + \xi_{it} \tag{A.5}$$

where the unobserved part of the equation is:

$$\xi_{it} \equiv (\lambda^Y(0)\theta_i + \epsilon_{it}^Y(0)) + D_{it}(\epsilon_{it}^Y(1) - \epsilon_{it}^Y(0)) + (\lambda^Y(1) - \lambda^Y(0))\theta_i$$

Thus, the consistency of OLS depends on whether the individuals know their unobserved latent ability endowment (θ_i) and act on it. In this case, net benefits of training ($I_{it}(h_t)$ in equation 1) depend on the unobserved latent ability endowment and the OLS estimator of π_1 (equation A.2) is inconsistent.

Since the inconsistency is originated because the analyst does not observe θ_i—and, thus, it cannot control for it—, one commonly-used approach is to assume that an individual fixed-effect factor drives selection bias. Even though the analyst does not observe θ_i, she can take advantage of the longitudinal nature of the data to eliminate this fixed effect. To see how, reorganize terms in equation (A.5) in the following way:

$$Y_{it} = \overbrace{\mu^Y(0)}^{\pi_0} + D_{it} \overbrace{(\mu^Y(1) - \mu^Y(0))}^{\pi_1} + \overbrace{[(\lambda^Y(1) - \lambda^Y(0))\theta_i + \lambda^Y(0)\theta_i]}^{u_i} + v_{it} \tag{A.6}$$

where $v_{it} \equiv \epsilon_{it}^Y(0) + D_{it}(\epsilon_{it}^Y(1) - \epsilon_{it}^Y(0))$, and note that the equation above is the standard fixed-effect regression. Here, the fixed effect u_i is a function of the unobserved productivity θ_i.

One way of estimating (A.6) is by taking First Differences (FD). Since we observe (Y_{it}, D_{it}) for various periods, we could run OLS on:

$$\Delta Y_{it} = \pi_1 \Delta D_{it} + \Delta v_{it},$$

where the fixed effect has been eliminated and the resulting error term is independent of D_{it}.[34] Therefore, by controlling for u_i, we can recover consistent estimates of π_1.

Which treatment parameter is the FD estimator recovering? Next, we show that the FD estimator identifies the average treatment effect (that is, $\pi = ATE$ as defined in equation A.1) only if the underlying model of counterfactual outcomes is independent of training histories—which means ignoring the dynamics we laid out in the previous section.

Consider the following assumptions:

Assumption 1. *In equation (6), $\mu^Y(h_t; 1) - \mu^Y(h_t; 0) = \pi_1$ for all $h_t \in \mathcal{H}_t$ and $t \in \mathcal{T}$.*

Assumption 2. *In equation (6), $\lambda^Y(h_t; 1) - \lambda^Y(h_t; 0) = 0$ for all $h_t \in \mathcal{H}_t$ and $t \in \mathcal{T}$.*

[34] When $T = 2$, the fixed-effect estimator is equivalent to the first-differences estimator. In this paper, we focus on the first-differences estimator, but the results are equivalent in the fixed-effect framework.

Assumptions 1 and 2 restrict the gains from training to be constant for all training histories. Assumption 2 rules out any potential gains to treatment for individuals with different levels of unobserved heterogeneity, thereby disregarding the possibility that individuals with higher levels of unobserved ability may enjoy larger returns to training. As a result, these assumptions not only impose strong restrictions within periods, but also across labor market and training histories; Assumptions 1 and 2 imply that the returns to training are equivalent for workers trained at time t with training histories $h \in \mathcal{H}_t$ and $h'' \in \mathcal{H}_t$ as well as for workers trained at time $t-1$ with histories $h' \in \mathcal{H}_{t-1}$. Furthermore, these assumptions imply absence of complementarities in the human capital accumulation process—a particularly strong restriction in the context of skills development in the labor market (Mincer, 1974).

Under these assumptions, we can show the following.

Proposition 1. *Suppose outcomes are determined by equation (6) and Assumptions 1 and 2. Then the FD estimator from equation (A.6) follows:*

$$\delta^{FD} = \pi_1 = \mu^Y(h_t;1) - \mu^Y(h_t;0) \qquad \textit{for } h_t \in \mathcal{H}_t, t \in \mathcal{T}$$

Proof. Let h and h' denote elements of $\mathcal{H}_t$ and $\mathcal{H}_{t-1}$. We can express the FD estimator as

$$\begin{aligned}\delta^{FD} &= 1/2 \times E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h)Y_{it}(h;1) - \sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')Y_{it-1}(h';0)\right] \\ &- 1/2 \times E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h)Y_{it}(h;0) - \sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')Y_{it-1}(h';1)\right]\end{aligned}$$

Given our assumption about counterfactual outcomes (equation 6), the equation above reduces to:

$$\begin{aligned}\delta^{FD} &= 1/2 \times E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h)(\mu^Y(h;1) + \lambda^Y(h;1)\theta_i) - \sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')(\mu^Y(h';0) + \lambda^Y(h';0)\theta_i)\right] \\ &- 1/2 \times E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h)(\mu^Y(h;0) + \lambda^Y(h;0)\theta_i) - \sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')(\mu^Y(h';1) + \lambda^Y(h';1)\theta_i)\right],\end{aligned}$$

and collecting terms, we have

$$\begin{aligned}\delta^{FD} &= 1/2 \times E\left[\sum_{h\in\mathcal{H}_t} H_{it}(h)(\mu^Y(h;1) - \mu^Y(h;0)) + \sum_{h\in\mathcal{H}_t} H_{it}(h)(\lambda^Y(h;1) - \lambda^Y(h;0))\theta_i\right] \\ &+ 1/2 \times E\left[\sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')(\mu^Y(h';1) - \mu^Y(h';0)) + \sum_{h'\in\mathcal{H}_{t-1}} H_{it-1}(h')(\lambda^Y(h';1) - \lambda^Y(h';0))\theta_i\right].\end{aligned}$$

Reducing the expression above by applying the expected value operator cannot yield ATE, because of two fundamental reasons. First, $H_t(h)$ is, in general, not independent of θ_i, since agents may sort into training at different periods based on their knowledge of θ_i. Second, even if $H_t(h)$ and θ_i were independent, the resulting weighted averages of treatment effects of t and $t-1$ may not necessarily have to be the same. Under assumptions 1 and 2, the second term in each square bracket collapses to 0 and the first term to a constant π_1. We have then

$$\delta^{FD} = 1/2 \times \pi_1 + 1/2 \times \pi_1 = \pi_1$$

As a result, assumptions 1 and 2 imply that the ATE equals π_1 across all training nodes and training histories (see equation A.1).

$\square$

Under Assumptions 1 and 2, Proposition 1 shows that the FD estimator recovers an average treatment effect which is constant in time and across histories. Hence, the FD recovers our parameter of interest only under the assumption of constant returns to training.

Another potential set of parameters of interest—specially relevant in the context of a dynamic setting—are treatment effects in a dynamic sense. Dynamic treatment effects can be of interest as they allow capturing potential complementarities in the returns to training. For instance, we may be interested in estimating the effect of training for a worker who has received training at time t and $t-1$ relative to a counter-factual history with no training in either period. Formally, this parameter can be defined as:

$$E[Y_{it}((h';1);1) - Y_{it}((h';0);0)], \qquad h' \in \mathcal{H}_{t-1} \tag{A.7}$$

Is the FD able to identify dynamic treatment effects as defined in equation (A.7)? One can show the FD estimator equals $E[\Delta Y_{it}|\Delta D_{it} = 1] - 1/2 \times E[\Delta Y_{it}|\Delta D_{it} = -1]$. Then, since the FD estimator requires using the sample of workers who have changed their participation decision in periods t and $t-1$, we cannot use the FD estimator to recover a dynamic treatment effect.

Table A.1 performs formal tests of Assumptions (1) and (2). Panel A presents the parameters associated with assumption (1). Assumption (1) requires that $\mu^Y(1) - \mu^Y(0) = \mu^Y(1,1) - \mu^Y(1,0) = \mu^Y(0,1) - \mu^Y(0,0) = \pi_1$ for $h \in \mathcal{H}_t$. In the implementation of the dynamic model, $\mu^Y(j, h_t)$ equals $X_i^Y \beta^Y(j, h_t)$ for $j \in \{0, 1\}$ and for all histories h_t. Using a F test, we test the null hypothesis that the three parameters are equal to each other. Our results indicate a strong rejection the null hypothesis (p-value < 0.01). Panel B presents the parameters associated with Assumption (2). In our context, this assumption requires $\lambda(1) - \lambda(0) = \lambda(1;1) - \lambda(1;0) = \lambda(0;1) - \lambda(0;0) = 0$ for $h \in \mathcal{H}_t$. This assumption implies that higher ability workers cannot enjoy additional returns to training across different time periods and training histories. We conduct the same F test and find the three parameters are statistically different from each other (p-value < 0.01). Therefore, we find evidence against the null hypothesis that fixed-effect estimators recover the ATE.

Table A.1: Testing Assumptions 1 and 2 for Validity of FD Estimators

(a) Assumption 1: $\mu^Y(h_t, 1) - \mu^Y(h_t, 0)$

		$\mu^Y(1) - \mu^Y(0)$	$\mu^Y(0,1) - \mu^Y(0,0)$	$\mu^Y(1,1) - \mu^Y(1,0)$
Coefficient		0.017	0.034	0.006
p-value of test	0.000			

(b) Assumption 2: $\lambda^Y(h_t, 1) - \lambda^Y(h_t, 0)$

		$\mu^Y(1) - \mu^Y(0)$	$\mu^Y(0;1) - \mu^Y(;,0)$	$\mu^Y(1;1) - \mu^Y(1;0)$
Coefficient		-0.003	0.007	-0.006
p-value of test	0.000			

Notes: We test for assumptions 1 and 2 using a simulated sample drawn from the estimated dynamic model. We show the p-value of the joint hypothesis of equality of parameters.

Table A.2: Estimated Returns to Different Job Training Histories

	Monthly Earnings After $t = 2$	
	(1)	(2)
Second-Period Trainees	0.209***	0.177***
	(0.008)	(0.007)
First-Period Trainees	0.0964***	0.0966***
	(0.009)	(0.007)
Trained in Both Periods	0.354***	0.266***
	(0.009)	(0.008)
Control Variables		✓
Observations	37089	37089
R^2	0.045	0.322

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors in parentheses. The omitted category is workers who do not participate in training in either year. Table A.2 presents the estimated returns to different training histories on average monthly earnings in the first quarter following second-year training. Column 1 does not include control variables. Column 2 includes baseline college entrance test scores, high school GPA, parental education and workers' gender as control variables.

B Estimated Parameters

Table B.1: Measurement system estimates

	Math PSU	Language PSU	High School GPA	Initial Salary
Constant	3.71	3.28	4.14	-2.58
	(0.14)	(0.14)	(0.14)	(0.06)
Male	0.13	-0.09	-0.38	0.11
	(0.01)	(0.01)	(0.01)	(0.01)
Age	-0.21	-0.18	-0.22	0.10
	(0.01)	(0.01)	(0.01)	(0.00)
Cognitive Factor	0.38	0.31	0.24	1.00
	(0.01)	(0.01)	(0.01)	(0.00)
Precision	1.20	1.12	1.14	3.58
	(0.01)	(0.01)	(0.01)	(0.03)
Observations	1000	1000	1000	1000

Note: The table displays the estimation results from the measurement system of test scores (equation 9). We obtain these estimates by simulating 1,000 values of parameters using our estimated posterior. The dependent variable are the standardized test score and the initial log earnings at the time of labor market entry. The earnings equation includes year-of-entry dummies. Standard errors are in parentheses. The loading on cognitive factor in the initial earnings equation is normalized to 1.

Table B.2: Training Probits: Estimated Parameters

	I	$I(0)$	$I(1)$
Constant	-1.32	-1.40	-0.43
	(0.04)	(0.05)	(0.08)
Gender	-0.03	-0.04	-0.09
	(0.02)	(0.02)	(0.03)
Firm average training hours	0.60	0.72	0.35
	(0.01)	(0.01)	(0.02)
Comuna average training hours	0.28	0.16	-0.03
	(0.07)	(0.08)	(0.13)
Cognitive Factor	0.11	0.15	0.27
	(0.01)	(0.01)	(0.02)
Precision	1.00	1.00	1.00
	(0.00)	(0.00)	(0.00)
Observations	1000	1000	1000

Note: We show the estimated parameters of the training probits (equation 5). We obtain these estimates by simulating 1,000 values of parameters using our estimated posterior. The dependent variable corresponds to the training dummy $I(h)$, for lagged training choice $h_t \in \{0, 1\}$. Standard errors are in parentheses.

Table B.3: Earnings Equations: Estimated Parameters

	$Y_i(h_2 = \{1\})$	$Y_i(h_2 = \{0\})$	$Y_i(h_3 = \{1,1\})$	$Y_i(h_3 = \{1,0\})$	$Y_i(h_3 = \{0,1\})$	$Y_i(h_3 = \{0,0\})$
Constant	5.91	4.69	5.82	5.13	4.60	4.88
	(0.17)	(0.33)	(0.23)	(0.38)	(0.50)	(0.49)
Gender	0.07	0.07	0.10	0.10	0.10	0.11
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Age	-0.03	0.08	-0.02	0.05	0.09	0.06
	(0.02)	(0.03)	(0.02)	(0.04)	(0.05)	(0.04)
Age^2	0.00	-0.00	0.00	0.00	-0.00	0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Cognitive Factor	0.59	0.58	0.55	0.56	0.56	0.55
	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Precision	114.82	87.79	17.68	19.60	18.36	20.85
	(4.95)	(4.78)	(0.20)	(0.44)	(0.49)	(0.60)
Observations	1000	1000	1000	1000	1000	1000

Note: We show the estimated parameters of the earnings process (equation 6). We obtain these estimates by simulating 1,000 values of parameters using our estimated posterior. The dependent variable corresponds to average monthly earnings $Y(h_t; j)$, for training choice $j \in \{0, 1\}$ and lagged training choice $h_t \in \{0, 1\}$. All earnings equations include year-of-entry dummies. Standard errors are in parentheses.

C Returns to Job Training: Stayer Sample

Figure C.1: Observed Training Choices for Stayers

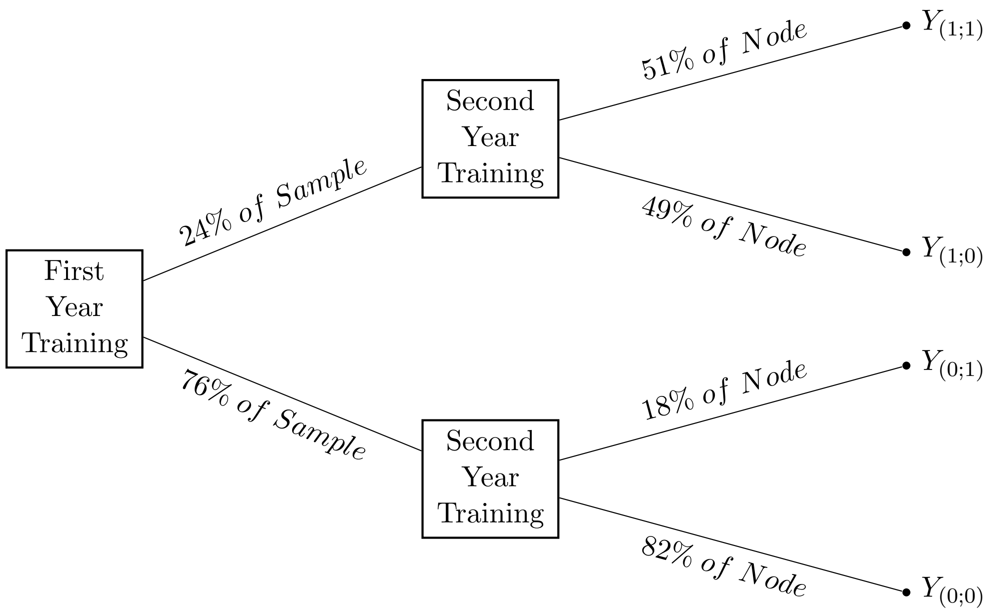

Notes: We present the decision tree through which workers decide whether to participate in training in each of their first two years in the labor force. In each node, we include the observed share of workers in our sample who decide to participate in training.

Table C.1: Summary Statistics: Stayer Sample

Variable	Mean	(Std. Dev.)
Male	0.44	(0.5)
Age at Graduation	17.81	(0.63)
Math PSU (Normalized)	-0.03	(0.98)
Verbal PSU (Normalized)	-0.01	(0.99)
High School GPA (Normalized)	0.04	(0.99)
Monthly Salary after First Year (USD)	567.01	(344.8)
Monthly Salary after Second Year (USD)	623.53	(392.41)
N	22247	

Note: The table presents summary statistics of the sample of stayers. The dependent variable is the monthly average of earnings in the first quarter following the training period. For simplicity, we refer to this variable as concurrent with the training decision. Tests scores (Math and Verbal) and high school GPA are standardized across the general population of test-takers to be of mean zero and variance 1.

Table C.2: Static Returns to Training: Stayer Sample

Treatment effect	$t=1$	$t=2$ $h_2=\{0\}$	$t=2$ $h_2=\{1\}$
ATE (percentage points)	2.54	2.95	0.35
	[2.51,2.58]	[2.87,3.04]	[0.18,0.52]
TT (percentage points)	2.60	3.02	0.07
	[2.53,2.68]	[2.82,3.22]	[-0.17,0.31]
$Pr(ATE<0)\times 100$	41.77	45.42	49.51
	[41.63,41.92]	[45.25,45.59]	[49.21,49.82]
$Pr(TT<0)\times 100$	41.65	45.14	49.96
	[41.34,41.95]	[44.75,45.53]	[49.51,50.41]

Notes: We report estimates of the Average Treatment Effects (ATE), Treatment on the Treated (TT), and likelihood of negative treatment effects across three training nodes. The first column show treatment effects after individuals make their first choice ($t=1$). The second and third columns show estimates for after individuals make a second choice ($t=2$), conditional on two possible choices in the first period ($D_1 \in \{0,1\}$). We show 95% confidence intervals in brackets.

Table C.3: Dynamic Returns to Training: Stayer Sample

	DATE	DTT	DTUT
Direct effect (short-term)	1.49	1.52	1.48
	(0.01)	(0.02)	(0.01)
	[37%]	[38%]	[37%]
Direct effect (medium-term)	2.61	2.63	2.60
	(0.02)	(0.04)	(0.02)
	[65%]	[66%]	[64%]
Continuation value	-0.07	-0.17	-0.04
	(0.02)	(0.04)	(0.02)
	[-2%]	[-4%]	[-1%]
Total	4.03	3.97	4.05
	(0.02)	(0.05)	(0.03)

Note: We estimate Dynamic Average Treatment Effects (DATE), Dynamic Treatment on the Treated (DTT) and Dynamic Treatment on the Untreated (DTUT) of training in $t = 1$ on the present value of earnings. Let $\widetilde{Y}_i(h_1; j) \equiv Y_i(h_1; j) + \rho\left(D_i(h_2)Y_i(h_2; 1) + (1 - D_i(h_2))Y_i(h_2; 0)\right) j \in \mathcal{H}_2 \equiv \{0, 1\}$ be the present value of earnings for training choice $j \in \{0, 1\}$. The individual-level dynamic treatment effect equals $\widetilde{Y}_i(1) - \widetilde{Y}_i(0)$, where

$$\widetilde{Y}_i(1) - \widetilde{Y}_i(0) = \underbrace{(Y_i(1) - Y_i(0))}_{\text{Direct effect (short-term)}} + \underbrace{\rho[Y_i(0;1) - Y_i(0;0)]}_{\text{Direct effect (medium-term)}} + \underbrace{\rho\left[D_{i2}(1)(Y_i(1;1) - Y_i(0;1)) - D_{i2}(0)(Y_i(1;0) - Y_i(0;0))\right]}_{\text{Continuation value}}.$$

We show the above decomposition for:

$$DATE \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0)\right],$$
$$DTT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 1\right],$$
$$DTUT \equiv E\left[\widetilde{Y}_i(h_1; 1) - \widetilde{Y}_i(h_1; 0) \mid D_i(h_1) = 0\right],$$

We present DATE, DTT, and DTUT as percentage of mean baseline of the present value of earnings ($E[\widetilde{Y}_i(0)]$). We present standard errors in parenthesis and the percentage contribution of each term in brackets.