

Dynamic Treatment Effects of Job Training*

Jorge Rodríguez¹, Fernando Saltiel², and Sergio Urzúa³

¹Universidad de los Andes

²McGill University

³University of Maryland and NBER

February 23, 2021

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, providing evidence of dynamic substitutability. Policy experiments illustrate how increasing the local availability of training programs may affect earnings heterogeneously across dynamic responses.

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-à-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

*Jorge Rodríguez, Universidad de los Andes, Chile; email: jrodriguez@uandes.cl; Fernando Saltiel, Department of Economics, McGill University; email: fernando.saltiel@mcgill.ca; 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.

¹Guvenen et al. (2020), Lise and Postel-Vinay (2020), Lise et al. (2016), and Saltiel et al. (2018) discuss skill (mis)match in the labor market. See Autor et al. (2003), Spitz-Oener (2006), Ingram and Neumann (2006), Acemoglu and Autor (2011), and Sanders and Taber (2012) for a literature on the changing returns to specific skills.

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.

We follow Heckman and Navarro (2007) and Heckman et al. (2016) to estimate a 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. We allow potential outcomes in period t to vary depending not only on current training choices but also on all possible past choices the individual has made up to period $t - 1$, thereby letting earnings to follow a flexible state-dependence process. As in Heckman et al. (2016), the joint distribution of counterfactual earnings across potential training histories is nonparametrically identified given a factor structure, in which a finite set of unobserved traits drives the correlation of unobservables across training choices and outcomes. By using a measurement system of pre-labor market test scores and pre-training wages, we are able to nonparametrically identify the distribution of two unobserved factors, which encompass workers’ latent ability and productivity. Exclusion restrictions are further exploited to identify policy-relevant parameters.

We implement our framework in the context of 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 2.5%. Meanwhile, the returns to second-period participation differ across first-period histories, reaching 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 4.4%. 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

²Formally, 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. $\frac{\partial^2 Y_t}{\partial D_{t-1} \partial D_t} \geq (\leq) 0$.

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. 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, 2019).

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 3.8%. 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, 2020; De Chaisemartin and d'Haultfoeuille, 2020; 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' unobserved heterogeneity 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 Vytlačil, 2001, 2005; Carneiro et al., 2010; Mogstad et al., 2018; Mogstad and Torgovitsky, 2018).

³A 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).

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 parameters 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 dynamic response types and presents evidence on the returns to the policy change. We conclude in Section 7.

2 Treatment Effect Framework

A Roy model framework helps us to characterize the dynamics of training decisions and labor market outcomes. It 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.


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 $t \in \mathcal{T} \equiv \{1, \dots, T\}$. We define \mathcal{H}_t as the set of all possible training decisions histories before 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.

⁴We 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.

Figure 1: Decision Tree



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.

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:

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

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.

We allow for counterfactual earnings to vary by training histories and current choices. The agent progresses through each node 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. Behind the definition of counterfactual earnings we implicitly assume that current outcomes do not vary by future choices. This assumption is referred as the no-anticipation condition and it implies that $Y_i(h_t; j)$ does not depend on choices at $t' > t$ (Abbring and van den Berg, 2003; Fruehwirth et al., 2016).⁵

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

⁵This assumption does not rule out a forward-looking behavior; agents can still make predictions about the future and act on them. However, conditional on a given information set (whatever this is at any given time), potential outcomes do not vary by future treatment choices.

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 $\tilde{Y}_i(h_t; j)$ be her present value of earnings associated with training option j given history h_t . Thus,

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

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 $\tilde{Y}_i(h_{t+1}; k)$ 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 $\tilde{Y}_i(h_t; 1) - \tilde{Y}_i(h_t; 0)$. To understand the different mechanisms through which job training affects long-term earnings, we follow Heckman et al. (2016) and decompose it as:

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

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 have 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:

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

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.⁶ 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 Background 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.⁷ 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.⁸ Training courses are held off-the-job, in centers managed by private providers.⁹ 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 soft skills; (ii) short-term degrees leading to specialization in specific disciplines; (iii) professional conferences.¹⁰ The predominant role played by the private sector in training course

⁶To relate to the definition of Heckman and Mosso (2014), we have dynamic complementarity (substitutability) if $\partial^2 \theta_{t+s+1} / \partial I_t \partial I_{t+s}$ is positive (non-positive), where θ_t is a type of skill and I_t is human capital investment. We follow this definition but adapt it to the case of discrete investments.

⁷In 2010, FT funded training courses for 920,688 workers, covering 12% of the labor force. Using data from Chile's national accounts, we estimate the total cost of FT to be 0.08% of GDP.

⁸FT subsidizes firms through a tax exemption, with a cap set at 1% of the firms' annual payroll. 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.

⁹There 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.

¹⁰Courses 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.

provision under FT resembles that of the Workforce Investment Act (WIA) in the United States.¹¹ While the workers targeted by the FT program are different than those in the WIA, 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 2010.¹² 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 as we do not 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.¹³ 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.¹⁴ 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

¹¹One of the main goals of the WIA, 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).

¹²Despite the heterogeneity in course-types discussed above, our data does not include information on the types of courses taken by workers. As such, we are unable to directly model the impact of different course types and intensities on labor market outcomes.

¹³This omission could influence the estimation of the distribution of unobserved heterogeneity. See Heckman (1981) and Arellano and Honore (2001) for a related discussion. We remark that given our focus on labor market entrants, our results may not generalize to an older sample of workers.

¹⁴A longer panel increases the number of counterfactual outcomes to identify and the associated parameters to estimate; the computational costs of estimating the effects of training increases exponentially with the number of periods considered. Thus, we chose to work with the minimum number of periods that enable us to identify dynamic treatment effects. As argued, even with this simple model, we can study rich dynamic effects, including dynamic complementarities and dynamic policy-relevant treatment effects (see Sections 5 and 6).

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.¹⁵ Furthermore, our choice of sample selection limits the scope of our results: by focusing on young workers starting their careers, our results on static and dynamic returns to job training might not generalize to human capital investments later in life. Overall, 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.

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)
Initial Monthly Salary (USD)	462.93	(277.8)
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 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 $h_3 = (h, h')$, 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 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.

¹⁵Our focus on the earnings dimension makes our paper comparable in scope to the on-the-job training literature.

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
Initial Monthly Salary (USD)	432.45	489.27	478.09	600.51
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
Average Hours Trained (First Year)	0	0	35.5	37.0
Average Hours Trained (Second Year)	0	31.6	0	40.1
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.

3.3 Reduced-Form Analysis

The descriptive statistics presented above highlight important differences in baseline characteristics and outcomes across training participation groups. We examine 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:

$$Y_{it} = \delta D_{it} + X'_{it}\beta + \alpha_t + \kappa_i + \varepsilon_{it}, \quad (4)$$

where Y_{it} represents the log of earnings for worker i in the first quarter following year t . 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.

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*** (0.007)	0.149*** (0.006)				
$D_{i2}(1)$			0.257*** (0.013)	0.161*** (0.011)		
$D_{i2}(0)$					0.209*** (0.008)	0.178*** (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*** (0.005)	0.179*** (0.004)	0.075*** (0.003)	0.009*** (0.003)
OLS	X			
OLS + controls	X			
OLS + initial wage	X			
Individual FE	X			
Observations ($N \times T$)	74178	74178	74178	74178

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} . Columns (1) and (2) uses all sample to estimate the regression of first-period earnings on first-period training (with and without control variables). Columns (3)-(6) present second-period earnings regressions onto second-period training conditioning the estimating samples on first-period training. 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) includes the same control variables along with the first monthly salary observed for each worker. Column (4) computes the first differences estimator. p -values are in parenthesis, where * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

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 workers' first monthly salary in the labor market, observed prior to training choices. The estimated returns fall to 7.5%. In the fourth column, we estimate a linear regression model with individual fixed effects. 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.

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’Haultfoeuille, 2020](#)). In the context of our model of choices and counterfactual outcomes, [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 [Appendix A](#), using our model estimates, we find evidence against the null hypothesis of lack of differential gains at conventional significance levels. Second, the parallel-trends assumption required for identification 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, 2020](#)), this estimator would not recover direct effect, continuation value and dynamic complementarity/substitutability parameters 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 Identification and Implementation

This section discusses the model specification, its implementation and identification. We then present the estimated parameters and analyze the extent of sorting on unobserved ability.

4.1 Model Specification

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 model the decision process using a linear-in-the-parameters index function:

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

where X_i^I is a vector of characteristics (observed by the econometrician and the agent) and $\eta_i^I(h_t)$ is an individual- and choice-level innovation (unobserved by the analyst). Critical for our analysis, this last term contains information used by the agent to make current choices (see next sub-section). 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:

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

where X_i^Y is a vector of observed characteristics and $\eta_i^Y(h_t; j)$ represents the error term, which is

unobserved by the econometrician but partially known to the agent. In this setup, the effects of unobserved 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 earnings vary accordingly. Furthermore, the earnings return to different paths of training choices—defined by h_t and j —vary across the latent heterogeneity distribution. Notice that $(\eta_i^Y(h_t; j), \eta_i^Y(h_{t'}; j'), \eta_i^I(h_t), \eta_i^I(h_{t'}))$ are correlated for any potential outcomes defined by j and j' , histories h_t and $h_{t'}$, and time periods t and t' . Heckman and Navarro (2007) discuss how exclusion restrictions can be used to secure the non-parametric identification of general dynamic discrete-choice models with similar complexities to ours.

4.2 Factor Structure

We start off the identification argument by imposing a factor structure on the error terms. In particular, we follow a well-established literature and posit that a low-dimensional set of latent variables (factors) generate dependence across choices and outcomes for all possible h_t (Aakvik et al., 2005; Carneiro et al., 2003; Heckman et al., 2006, 2016, 2018). Thus, we assume:

$$\eta_i^I(h_t) = \boldsymbol{\theta}_i \boldsymbol{\lambda}^I(h_t) + \varepsilon_i^I(h_t), \quad (7)$$

$$\eta_i^Y(h_t; j) = \boldsymbol{\theta}_i \boldsymbol{\lambda}^Y(h_t; j) + \varepsilon_i^Y(h_t; j). \quad (8)$$

where $E(\boldsymbol{\theta}_i) = 0$ and $\boldsymbol{\theta}_i$ represents a finite set of fixed, latent endowments known by the agent but not the econometrician, which are independent of observed individual characteristics. $\varepsilon_i^I(h_t)$ and $\varepsilon_i^Y(h_t; j)$ are idiosyncratic shocks to the decision process and outcomes, respectively, that the agent cannot anticipate. As such, these are independent of observed characteristics and $\boldsymbol{\theta}_i$, and independent across choices and outcomes as well as across time and training histories.¹⁶ In this way, conditional on X_i^Y and X_i^I , $\boldsymbol{\theta}_i$ generates all cross-correlations of outcomes and choices across histories. Intuitively, the correlation across unobservables arises only through those components the individual knows (and thus can act upon them) but not from those who cannot anticipate. As pointed out in the literature, this setting extends the matching-on-observables assumption by matching on unobserved heterogeneity.¹⁷

Measurement System. We rely on pre-training individual-level human capital measures to non-parametrically identify the distribution of $\boldsymbol{\theta}$ as well as the parameters governing the outcome equations. Let $\boldsymbol{\theta} \equiv [\theta_1, \theta_2]$, so we assume a two-factor model. We interpret θ_1 as latent ability and θ_2 as unobserved labor market productivity. In what follows, we configure the measurement system which is the basis of our identification strategy.¹⁸

First, we consider a set of K pre-labor market achievement test scores, T_{ik} with $k = \{1, \dots, K\}$. In our application, these correspond to the results in college admission tests and high school GPA,

¹⁶Formally, 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 $(h'; j')$.

¹⁷Abbring and van den Berg (2003) propose a related treatment effects framework and provide identification results. Our application differs substantially from their framework, as they consider a continuous time model where outcome variables are durations and treatment effects under one possible treatment. We consider participation decisions across multiple periods. See also Abbring and Heckman (2007).

¹⁸This strategy has been applied in various papers on human capital investment decisions. Carneiro et al. (2003), Hansen et al. (2004), Heckman et al. (2006), Agostinelli and Wiswall (2016a), Agostinelli and Wiswall (2016b) and Attanasio et al. (2020) are just a few examples.

all observed before the individual enters the labor market. For each of them, we write:

$$T_{ik} = X_{ik}^T \beta_k^T + \theta_{i1} \lambda_k^T + \varepsilon_{ik}^T, \quad (9)$$

where X_{ik}^T is a vector of exogenous control variables, θ_{i1} captures latent ability and ε_{ik}^T is assumed to be independent across test scores $k \in \mathcal{K}$ and of all other error terms in the model.¹⁹

Given our interest in analyzing the returns to job training for employed workers, we consider an additional dimension of unobserved heterogeneity: worker’s latent labor market productivity. In order to avoid any reverse causality from training to productivity, which would jeopardize its identification, we rely on information about workers’ labor market outcomes prior to any training choices.²⁰ In particular, we take advantage of information on workers’ monthly earnings, W_{il} with $l = \{1, \dots, \mathcal{L}\}$, during that period, which we model as:

$$W_{il} = X_{il}^W \beta_l^W + \theta_{i1} \lambda_{l1}^W + \theta_{i2} \lambda_{l2}^W + \varepsilon_{il}^W, \quad (10)$$

where X_{il}^W includes control variables, θ_2 captures latent labor market productivity and ε_{il}^W is assumed to be independent of all other error terms in the model. Latent ability is also allowed to determine earnings.

Equations (9) and (10) conform a triangular measurement system. Within this framework, Carneiro et al. (2003) and Hansen et al. (2004) show that non-parametrically identifying the distributions of the latent ability factor ($f_{\theta_1}(\cdot)$) and error terms ($f_{\varepsilon_{ik}^T}(\cdot)$) along with the slopes (λ_k^T), requires at least three test score measures ($\mathcal{K} \geq 3$). On the other hand, the non-parametric identification of the distributions $f_{\theta_2}(\cdot)$ and $f_{\varepsilon_{il}^W}(\cdot)$ along with the slopes of the earnings equations requires $\mathcal{L} \geq 2$. Since both requirements are met in our context, we can identify the slopes and the distributions of the latent factors and error terms (up to two normalizations). The formal identification argument is presented in Appendix B.²¹ Since the distribution of θ is identified non-parametrically, we can impose a flexible parametrization for its estimation. We return to this issue below.

4.3 Exclusion Restrictions

Exclusion restrictions (instruments) represent the second building block for securing the identification of the treatment parameters of interest. In the literature, these are typically deployed to address endogeneity concerns. In our case, however, latent factors account for the potential correlation between unobservables across the choice and outcome equations. Alternatively, instruments can be used to relax parametric assumptions commonly imposed in switching regression models.²² Despite the fact our setting fits into this class, data limitations and the complexity of the model tilt the scale in favor of our linear-in-the-parameters specification.

¹⁹Notice that covariates are independent of θ_1 . This assumption underlines our interpretation of θ_1 . For instance, if we control for mother’s education in the test scores equations, then θ_1 would capture the unobserved ability orthogonal to that variable. Urzua (2008) and Attanasio et al. (2020) relax this structure.

²⁰While the literature on latent factors has traditionally relied on test score measures to identify unobserved ability, it largely focuses on human capital investment in the context of schooling choices. Our context is different, as we examine post-schooling training choices. We thus extend the literature by identifying baseline (or pre-training) latent labor market productivity using initial wages. We thank two anonymous referees for this suggestion.

²¹Ashworth et al. (2017) present a similar argument in the context of identifying the returns to schooling.

²²Consider, for instance, a more general version of the index function shaping training decisions (equation (5)): $I_i(h_t) = \beta^I(X_i^I) + \eta_i^I(h_t)$ with $h_t \in \mathcal{H}_t$ and $t \in \mathcal{T}$, where $\beta^I(\cdot)$ is an unknown function. Following the arguments in Matzkin (1994), assumptions on the instruments (existence and support) pave the way to the non-parametric identification of both $\beta^I(\cdot)$ and the distribution of the error term $\eta_i^I(h_t)$. See Heckman and Navarro (2007) for a detailed discussion.

Instead, in this paper we exploit a more modest but not less important characteristic of exclusion restrictions. On conceptual grounds, instruments encompass observed characteristics which affect how individuals dynamically sort across training status. Consequently, since individual’s responses to counterfactual scenarios can depend on observed and unobserved characteristics, we must articulate the specification of decision rules.

Consider, for instance, the identification of treatment effects associated with changes in the training choice model in period t . Contemporaneous and future compliers may possess different observed and unobserved characteristics relative to the average individual, so the resulting outcomes might not be identified by sample averages. In fact, this explains why reduced form models cannot identify local/marginal treatment effects without further structure (Heckman and Urzua, 2010). Thus, we shall articulate the specification of decision rules across potential histories and time periods, including any variable that might affect $I_i(h_t)$ but not the outcomes. This, along with the parameters and distributions already identified, enables the identification of counterfactuals that depend on past or present training decisions which, in turn, can be influenced by exogenous shocks (Mogstad and Torgovitsky, 2018). This secures the identification of the treatment parameters of interest (Heckman and Vytlacil, 2001).

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.²³ 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 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, the results in our Appendix C shows that we do not find evidence of a large correlation between our (residualized) instruments and baseline variables not used in estimation (Figure C.1). Moreover, to determine their empirical importance, we estimate the model without exclusion restrictions. We present the results in Table C.1. Since the estimated returns to first-period training differ vis-a-vis those in our preferred specification, we note that exclusions are indeed quantitatively important as determinant of sorting across decisions and, consequently, of treatment effects.

4.4 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:

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

where $f_{T_k}(\cdot)$ is the conditional density function of test score k , $f_{W_l}(\cdot)$ is the conditional density function of initial wage l , $F(\boldsymbol{\theta})$ represents the cumulative distribution function of the latent factors,

²³Related 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).

and

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

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).²⁴ However, for estimation purposes, we do adopt a flexible functional form for each component ($n \in \{1, 2\}$) of $\boldsymbol{\theta}$:

$$\theta_n \sim \rho_{1,n} N(\tau_{1,n}, \sigma_{1,n}^2) + \rho_{2,n} N(\tau_{2,n}, \sigma_{2,n}^2),$$

that is, a mixture of two normal distributions with means $(\tau_{1,n}, \tau_{2,n})$, probabilities $(\rho_{1,n}, \rho_{2,n})$ with $\rho_{1,n} + \rho_{2,n} = 1$ and variances $(\sigma_{1,n}^2, \sigma_{2,n}^2)$. We estimate the model by Markov Chain Monte Carlo (MCMC) and thus inference follows standard Bayesian arguments.²⁵ We present our estimated parameters generating 500 draws from the estimated posterior and report the resulting average. Using the Bernstein-von Mises theorem, we obtain the associated standard errors exploiting the standard deviations computed from these draws. A similar procedure is used to compute functions of the structural parameters (for example, ATE, TT, Dynamic ATEs, etc) and the associated standard errors.

Table 4: Variables Used in Implementation of the Model

Variables	Earnings Equation	Training Probit	Test Score Eq.	Initial Wages
Constant	Yes	Yes	Yes	Yes
Gender	Yes	Yes	Yes	Yes
Age at Test			Yes	
Age at Entry				Yes
HH Size			Yes	
Parents' Education			Yes	
Parents' Employment			Yes	
Age in Year t	Yes	Yes		
Latent Ability	Yes	Yes	Yes	Yes
Latent Productivity	Yes	Yes		Yes
Average Training Hours at Firm		Yes		
Average <i>Comuna</i> 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.

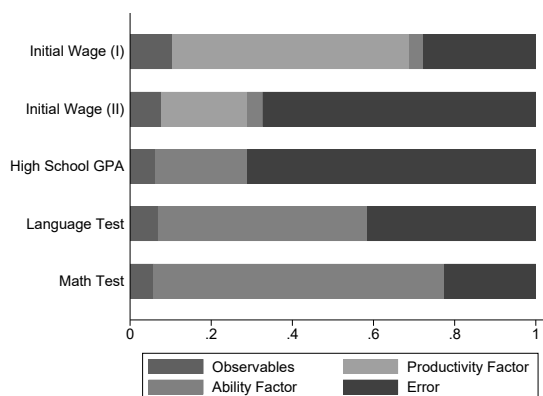
²⁴Even though we assume normal disturbances, note that our identification argument does not rely on normality and we only assume it for computational convenience. We have alternatively estimated a version of the model in which the error terms in the test score equations follow a mixture of normal distribution. The estimated results are in line with those presented in the main text and available upon request. We additionally assume that $\theta_1 \perp \theta_2$, yet remark this assumption can be relaxed as in Heckman et al. (2018).

²⁵We estimate the model using a Gibbs sampler as the MCMC algorithm. This delivers a sequence of realizations that can be used to characterize the distributions of the parameters governing the model. Formally, if $L(Y | \Psi)$ denote the likelihood of the function for a given set of parameters Ψ , then the posterior distribution, $g(\Psi | Y)$, is such that $g(\Psi | Y) \propto L(Y | \Psi)g(\Psi)$, where \propto stands for “proportional to”. We use the Gibbs sampling as a way of obtaining the posterior distribution of the parameters. As it is standard, at each iteration of the algorithm we form a conditional distribution of a set of parameters, ψ_j , given by $g(\psi_j | \psi_{-j}, Y)$, where ψ_{-j} is the sub-vector of all structural parameters besides the sub-set ψ_j . The sample mean obtained from repeated draws of $g(\psi_j | \psi_{-j}, Y)$ is then used for updating the conditional distribution of $g(\psi_k | \psi_{-k}, Y)$. We repeat this process to update all parameters, across iterations until convergence. Appendix D describes the algorithm in detail.

Table 4 shows the variables we include in the measurement system, training probit and earnings equation. We consider two PSU test scores (math and language) along with students’ high school GPA, which we allow to depend on the age at the time of PSU, along with family background information. The initial wage measures depend on the age at labor market entry and gender as well as the two latent factors. 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.

Estimation Results. In Appendix E, we show the estimated parameters of our econometric model. In Table E.2, we present the estimated parameters of the test score measures. The latent ability factor loads positively on both test score measures and high school GPA. Meanwhile, both factors load positively on the initial wage measures, yet the coefficient is larger on the latent productivity factor. To understand the relative contribution of observed characteristics and latent ability vector to test score measures, we present a variance decomposition in Figure 2. Latent ability explains 71%, 51% and 22% of the variance in the math PSU and verbal PSU scores and high school GPA, respectively. Meanwhile, latent productivity explains 50-56% of the variance in the initial wages, while the ability factor explains just 2-3% of the corresponding variance. In Figure 3, we depict the distribution of unobserved heterogeneity. The latent ability component (θ_1) deviates significantly from normality, whereas θ_2 does not. We thus remark the importance of allowing for flexibility in the estimation of the θ distribution.

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.

Tables E.3 and E.4 present the estimated parameters of the training and earnings equations. Across all choice nodes, women are more likely to participate in training and individuals with higher latent ability and productivity 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.²⁶ The earnings equations indicate that males outearn women by upwards of 0.07 log points.

²⁶The only statistically insignificant coefficient corresponds to the one on local-level course availability in the

No discernible pattern arises with respect to the age-earnings profile. Both components positively impact earnings, and this relationship holds across all training nodes.

Figure 3: Distribution of Latent Factors



Note: In Figure 3, we show the estimated density of the latent factors. 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. The first panel presents the density of the latent ability factor (θ_1), whereas the second panel shows the corresponding density for the latent productivity factor (θ_2).

Lastly, we assess the model’s accuracy in predicting observed outcomes and choices. Figure E.1 compares observed and simulated training histories. The model matches training decisions well, both in the first and second year. Table E.1, on the other hand, contrasts and formally tests for the equality of data versus model-simulated first moments from the (log) earnings distributions. Panels A and B display the results for $t = 1$ and $t = 2$, respectively. Our empirical strategy precisely reproduces the overall (unconditional) moments in both periods, so we cannot reject the null hypotheses of equality between actual and model (last column).²⁷ Conditional on training choices, the model also provides a good approximation of the first moments – the objects of interest for this paper – although, in general, formal tests indicate statistical differences between them. In fact, the model predicts earnings within only 0.05 log points of the observed means across all nodes but $h_2 = \{1\}$ and $j = 1$. The difficulties of matching this moment are expected as only 10% of the sample reach this node. The model, we conclude, mimics the key features of the data.

4.5 Selection on Unobserved Characteristics

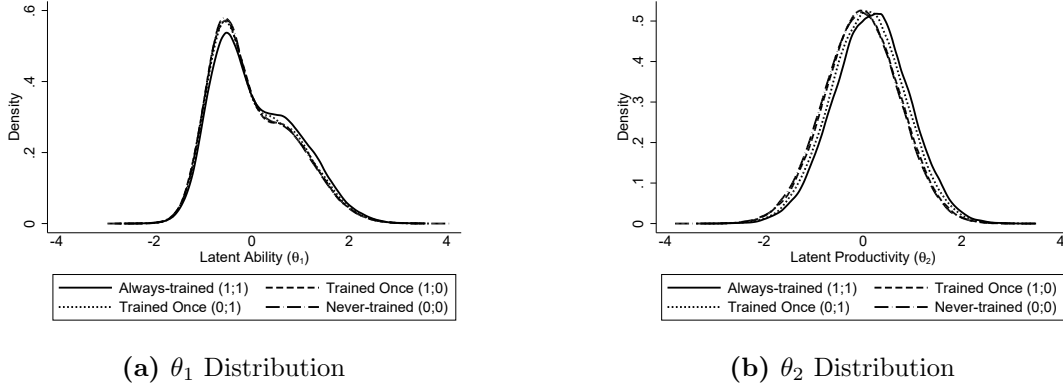
Figure 4 compares the density of the latent factors for workers choosing different training paths. Training choices are denoted by $(h_2; j)$, where h_2 represents the training history prior to $t = 2$ (or the first period decision) and j captures the training choice at $t = 2$. We find significant differences in latent ability distribution across one-time participants, depending on the timing of the decision. We find that for those who choose training in the first period but not in the second, the distribution of both θ_1 and θ_2 almost entirely overlaps with the density of the never-trained group. On the other hand, the distribution of the latent factor for those only trained in the second period clearly surpasses the never-trained group. In both the ability and productivity components, the latent skills of the always-trained group (1, 1) surpass the never-trained group by 0.2-0.3 standard deviations, fitting in with the test score differences shown in baseline test scores in Table 2. Overall, we find

training equation of period $t = 2$, conditional on training in the first period.

²⁷Following the conventional switching regression model, we define the overall (unconditional) moment as $Y_t = Y(h_t, 1) \times D(h_t) + Y(h_t, 0) \times (1 - D(h_t))$.

evidence of sorting on unobservables, as higher-skilled workers are more likely to have participated in job training.

Figure 4: Distribution of Latent Factors by Training History



Note: Figure 4 shows the estimated density of the latent factors 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. The first panel presents the density of the latent ability factor (θ_1), whereas the second panel shows the corresponding density for the latent productivity factor (θ_2).

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 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 distribution of unobserved heterogeneity.

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 $E[\cdot]$ 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,²⁸

$$ATE(h_t) \equiv E[Y_i(h_t; 1) - Y_i(h_t; 0)]. \quad (11)$$

²⁸Formally, the exact definition is as follows:

$$ATE(h_t) \equiv \int \dots \int (Y_i(h_t; 1) - Y_i(h_t; 0)) dF(\mathbf{X}, \theta, \varepsilon),$$

where $F(\mathbf{X}, \theta, \varepsilon)$ is the cumulative joint distribution of observables and unobservables. Under the independence assumption of our setting, $f(\mathbf{X}, \theta, \varepsilon) = f(\mathbf{X})f(\theta)F(\varepsilon)$.

We can additionally define 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:

$$TT(h_t) \equiv E[Y_i(h_t; 1) - Y_i(h_t; 0) \mid D_i(h_t) = 1]. \quad (12)$$

We present the estimated static returns to job training in Table 5. The average short-term returns to first-period training ($ATE(h_1)$) equal 2.5%, 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.^{29,30} For second-period training, conditional on not training in the first year ($h_2 = \{0\}$), we find larger static returns, reaching 4.0%. This effect significantly exceeds that for second-year participation for workers who had been trained in the first period ($h_2 = \{1\}$), which equals 0.8%. 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 and TT parameters are positive on average, first-period training lowers earnings for 43.5% of treated workers.³¹ 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. However, negative returns to training could arise from the comparison of two regimes. One in which training leads to the accumulation of human capital irrelevant to the firm, and one in which labor market experience (no training) leads to firm-specific human capital. The comparison of outcomes across these states could result in negative returns to training.

Table 5: Static Returns to Job Training (in %)

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
ATE (percentage points)	2.54 [2.50,2.58]	3.99 [3.91,4.08]	0.78 [0.62,0.94]
TT (percentage points)	2.58 [2.50,2.66]	4.07 [3.84,4.30]	0.81 [0.69,0.93]
$Pr(ATE < 0) \times 100$	43.65 [43.54,43.77]	45.15 [45.03,45.28]	49.02 [48.77,49.27]
$Pr(TT < 0) \times 100$	43.57 [43.32,43.81]	45.07 [44.78,45.37]	49.03 [48.65,49.41]

Notes: Table 5 presents the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters along with the share of workers who experience 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.

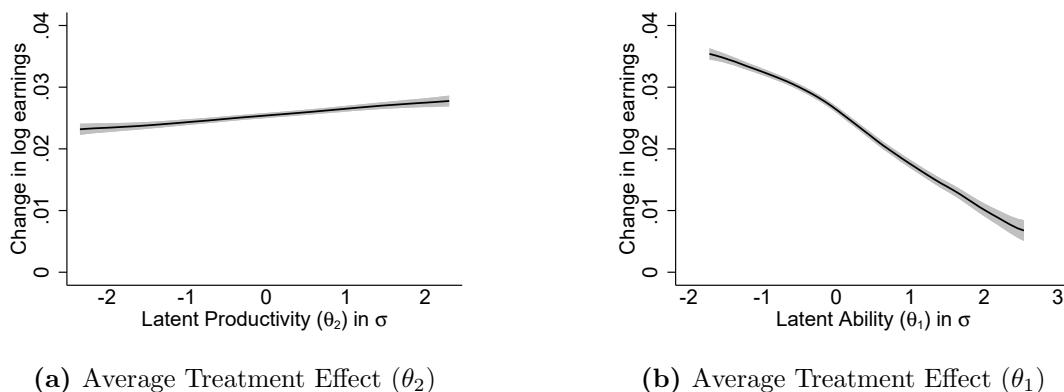
²⁹Appendix F documents that selection bias on unobservable characteristics plays a critical role in explaining the differences between reduced form estimates and the treatment effects recovered using estimated model parameters.

³⁰Appendix G presents evidence from an alternative specification which includes a full set of observables in the wage equations, showing similar static and dynamic returns to training.

³¹Table G.3 calculates the share of workers with negative returns *without* considering variation from the idiosyncratic wage shocks. This alternative interpretation indicates that 15% of workers earn negative returns from training.

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, 2020; De Chaisemartin and d’Haultfoeuille, 2020; Goodman-Bacon, 2018).

Figure 5: Heterogeneous Returns to First-Period Training Participation

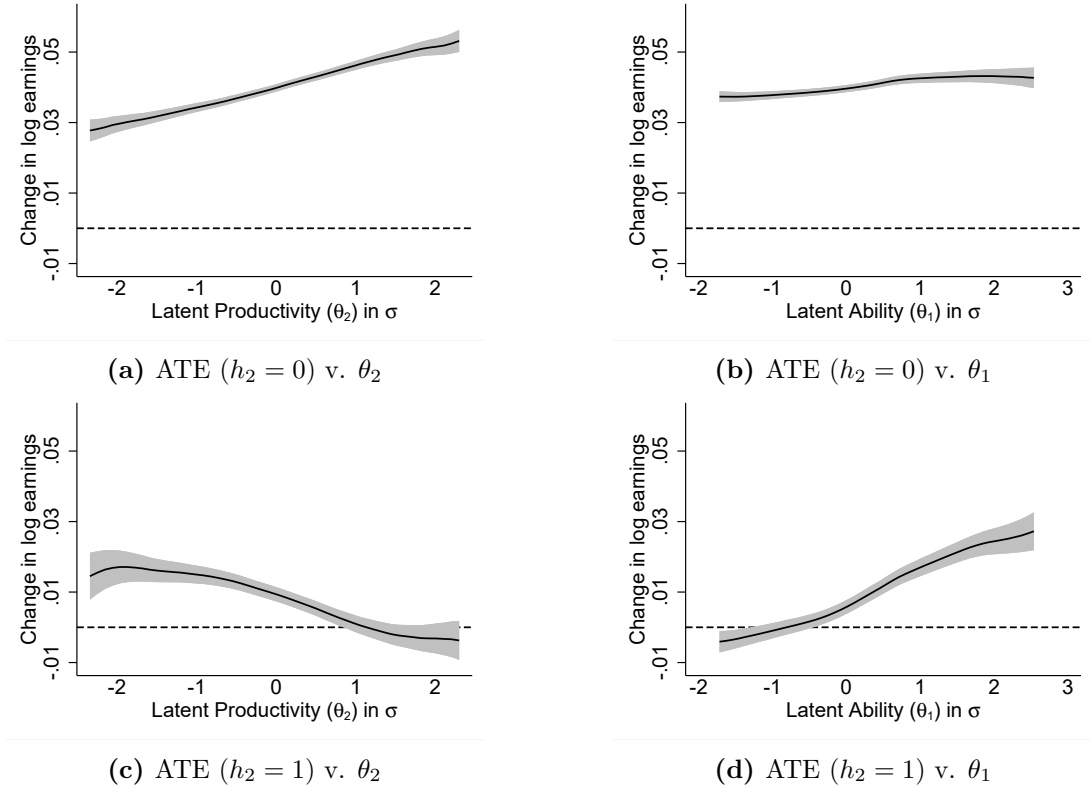


Note: In Figure 5, we estimate local polynomial regressions of the estimated ATE parameter for the first training event (at $t = 1$) against the distribution of latent ability (θ_1) and latent productivity (θ_2).

Figures 5 and 6 examine heterogeneous effects of training across the two latent factors.³² In the first year of training (Figure 5), we find largely flat returns to training for workers across the latent productivity distribution (θ_2), yet find sizable heterogeneity in the latent ability component, as less-skilled individuals experience far larger returns vis-a-vis their high-skilled counterparts. The returns to second year training (Figure 6) show that among workers who had not been initially trained ($h_2 = \{0\}$), program participation has larger effects on earnings for high latent productivity workers, exceeding 5% for those in the top decile of the skill distribution, while remaining close to 3% for individuals in the bottom decile. On the other hand, we do not find significant heterogeneity in the latent ability distribution. Lastly, for first-period participants ($h_2 = \{1\}$), the impact of second-period training for workers above the median of the latent productivity distribution (θ_2) is not different from zero, while exceeding 2% for those in the top of the latent ability distribution. The difference in heterogeneous returns across the two latent factors remarks the importance of considering multiple dimensions of unobserved heterogeneity when analyzing the returns to job training participation.

³²To compute these figures, we simulate outcomes and choices drawing different values from the distribution of θ . 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.

Figure 6: Heterogeneous Returns to Second-Period Training Participation



Note: In Figure 6, we estimate local polynomial regressions of the estimated average treatment effect of second-period job training participation against the distribution of latent ability (θ_1) and latent productivity (θ_2), conditional on first-period choices ($h_2 = j$ where $j \in \{0,1\}$). The top 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 against the latent productivity distribution (θ_2).

The estimated static treatment effects reveal heterogeneous impacts across different decision margins, job training histories, and latent factors. As such, the constant effect framework required in fixed-effect estimators is rejected in favor of a model of differential returns from job training.³³

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,

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

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

³³In Appendix H, we further examine the intensive-margin returns to job training. Table H.1 presents reduced form evidence showing larger returns to courses lasting longer than 20 hours. To recover the returns to longer courses within our model, we restrict the sample to workers who either participated in courses lasting longer than 20 hours or did not take-up training at all. Tables H.2 and H.3 show larger static and dynamic returns to participation in longer training courses vis-à-vis those presented in the main model, respectively. As such, future work should account for the effects in the intensity of training courses.

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 $t = 2$.

Table 6: Dynamic Returns to First-Period Job Training (in %)

	DATE	DTT	DTUT
Direct effect (short-term)	1.59 (0.01) [36%]	1.61 (0.02) [37%]	1.59 (0.01) [36%]
Direct effect (medium-term)	2.82 (0.02) [64%]	2.80 (0.05) [64%]	2.83 (0.02) [64%]
Continuation value	0.01 (0.02) [0%]	-0.04 (0.04) [-1%]	0.02 (0.02) [0%]
Total	4.42 (0.02)	4.38 (0.05)	4.43 (0.03)

Notes: Table 6 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. We present DATE, DTT, and DTUT as percentage of mean baseline of the present value of earnings ($E[\tilde{Y}_i(0)]$). We present standard errors in parenthesis and the percentage contribution of each term in brackets.

In Table 6, we present estimates from the following dynamic treatment effect parameters:

$$DATE \equiv E \left[\tilde{Y}_i(h_1; 1) - \tilde{Y}_i(h_1; 0) \right], \quad (13)$$

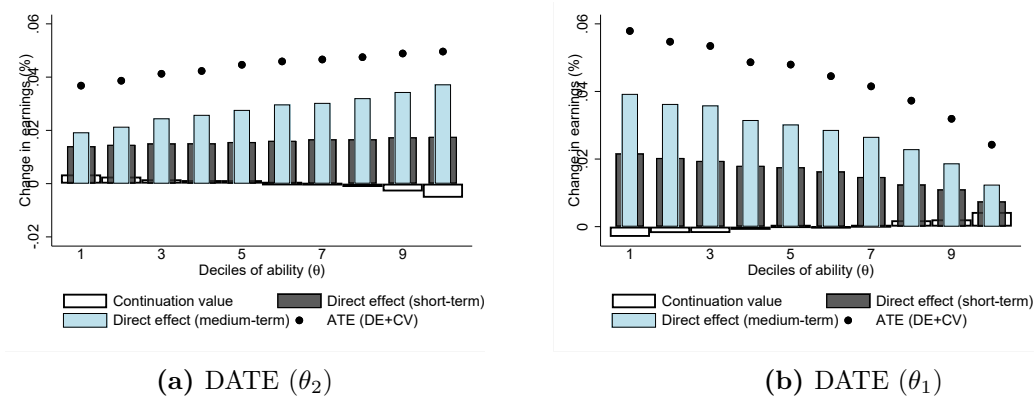
$$DTT \equiv E \left[\tilde{Y}_i(h_1; 1) - \tilde{Y}_i(h_1; 0) \mid D_i(h_1) = 1 \right], \quad (14)$$

$$DTUT \equiv E \left[\tilde{Y}_i(h_1; 1) - \tilde{Y}_i(h_1; 0) \mid D_i(h_1) = 0 \right], \quad (15)$$

and decompose them into short- and medium-term direct effects and continuation values.³⁴ 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 4.4%. This result is driven largely by the direct effect of job training, as first-period participation increases earnings (in present-value terms) by 1.6 and 2.8% one and two years after training, respectively. However, early training increases the returns to training in the second period—the continuation value—by 0.01%, though the effect is not statistically significant. This result is consistent with the estimates of Table 5, 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 explain the entirety of the estimated *DATE* (100%). As with the static treatment effects, we find similar effects in the dynamic *TT* and *TUT* parameters.³⁵

³⁴For this exercise, we assume a discount factor $\rho = 1/(1.05)$. In addition, we present our results as % increase

Figure 7: Heterogeneous Dynamic Returns to First-Period Training Participation



Note: In Figure 7, we estimate Dynamic Average Treatment Effects (DATE) of training in $t = 1$ on the present value of earnings across deciles of latent ability (θ_1) and latent productivity (θ_2) as defined in Section 5.

In Figure 7, we examine heterogeneous dynamic treatment effects across the two latent factors, and decompose them into direct effects and the continuation value. The first panel examines heterogeneity in the latent productivity component. Confirming the results shown in Figure 5, the short-term direct effects of training are largely flat across the θ_2 distribution. In contrast, the medium-term direct effects are larger for more productive workers. Lastly, we find that the continuation value of training is largely flat, leading to a DATE which does not differ significantly across the latent productivity dimension. Nonetheless, we find significant heterogeneity across workers' latent abilities. Both the short- and medium-term direct effect components are decreasing across the θ_1 distribution, while the continuation value of training is largely flat. We find sizable differences in the DATE of job training, as workers in the bottom latent ability decile experience returns reaching close to 6%, yet the returns for their top θ_1 decile peers barely exceed 2%.

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.³⁶ On the other hand, we find negative continuation values from job training for high productivity 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 val-

from the average baseline present value of earnings.

³⁵While the distribution of θ is identified non-parametrically, its distribution is estimated through a mixture of two normals. In Appendix I, we examine the robustness to estimating the latent factors using a mixture of three normal distributions. The results are robust, as the first period returns to training reach 2.8% and the dynamic ATE equals 4.6%, largely fitting in with the estimated returns presented in this section.

³⁶We further consider wage outcomes for workers two years after the last training event, allowing us to define 'long-term' earnings for worker i choosing training option j given her training history h_2 : $\widehat{Y}_i(h_2; j) = Y_{i,t+1}(h_2; j) + \rho \times Y_{i,t+2}(h_2; j)$. We re-estimate the model using the 'long-term' earnings measure as the earnings outcome following the second training decision. We present the static and dynamic returns to job training in Tables H.4 and H.5, respectively. We find larger returns to second-period participation along with larger dynamic returns to training. As such, we find a small, yet positive continuation value of training (5%). The estimated continuation value is still significantly lower from those found by Heckman et al. (2016), remarking the differences in the returns to human capital accumulation during and post-schooling.

ues and dynamic complementarity (substitutability) in the two-period model can be expressed as follows:

$$\begin{aligned}
 & \underbrace{E [D_i(1)(Y_i(1;1) - Y_i(1;0)) - D_i(0)(Y_i(0;1) - Y_i(0;0))]}_{\text{Continuation Value}} = \\
 & \quad \underbrace{E (Y_i(1;1) - Y_i(1;0)) - (Y_i(0;1) - Y_i(0;0))}_{\text{Dynamic Complementarity/Substitutability}} + \\
 & \underbrace{E (D_i(1) - 1)(Y_i(1;1) - Y_i(1;0)) - (D_i(0) - 1)(Y_i(0;1) - Y_i(0;0))}_{\text{Dynamic Sorting Gains}} \tag{16}
 \end{aligned}$$

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: $E [Y_i(1;1) - Y_i(1;0)] > E [Y_i(0;1) - Y_i(0;0)]$.

In Figure 8, we present evidence from a local polynomial regression of dynamic complementarity (substitutability) parameter onto the latent productivity distribution (equation (16)). First-period job training lowers the return from subsequent participation, independent of second-period decisions: on average, we find evidence of dynamic substitutability (-2.6%). On the other hand, dynamic sorting gains are positive (2.6% on average) and increasing across the latent productivity distribution, exceeding 3.5% for those in the top decile, fitting in with the sorting patterns documented in Section 4.5.³⁷

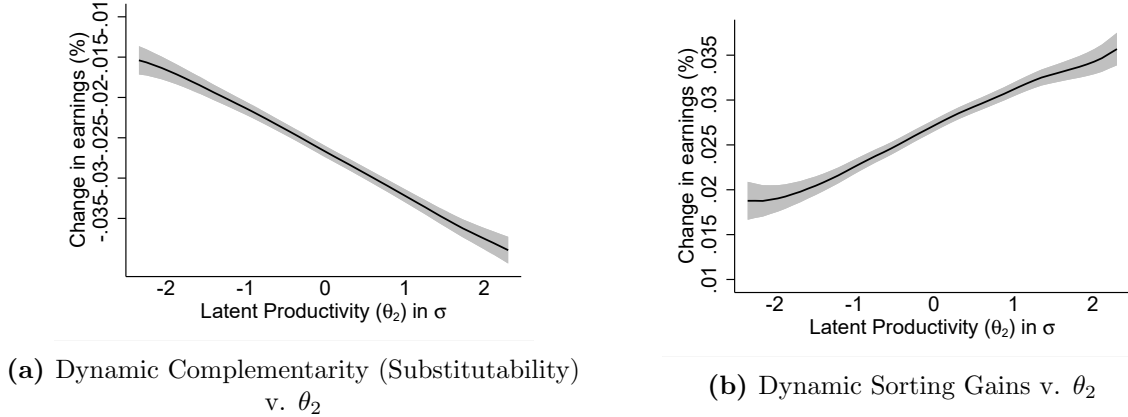
We note that dynamic substitutability, in the context of job training, may arise for various reasons. First, job training could comprise multiple courses covering topics in the same area, with workers starting in a baseline course and subsequently taking part in more complex coursework.³⁸ In this setting, 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.³⁹ Workers could rationally follow such a strategy as the returns to the high-payoff courses could be enjoyed over a longer time horizon.

³⁷We note that the heterogeneity in the dynamic complementarity (substitutability) parameter is largely muted across the latent ability distribution.

³⁸For example, workers first need to learn to operate a computer before taking a course on a specific software.

³⁹In this context, workers would first take courses directly related to their industry or occupation and subsequently participate in foreign language courses, for example.

Figure 8: Heterogeneous Dynamic Complementarity (Substitutability) and Dynamic Sorting



Note: In Figure 8, 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 latent productivity (θ_2).

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 who do not switch firms in their first two years in the labor force (“stayers”) in Appendix J.⁴⁰ As we find similar static and dynamic returns to training in the stayer subsample (Tables J.1 and J.2), 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 and heterogeneous returns to job training participation. 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.

⁴⁰This 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.

6.1 Policy Intervention

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.⁴¹ Even though the policy change only affects 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 $\{D_i^{a'}(h_1) = 1, D_i^a(h_1) = 0\}$.⁴² 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: $Y_i^a(h_t; j) = Y_i^{a'}(h_t; j) = Y_i(h_t; j)$, $j \in \{0, 1\}$.⁴³

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. 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.⁴⁴

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 policy, in $t = 1$:

$$Y_{i,1}^a \equiv D_i^a(h_1)Y_i(h_1; 1) + (1 - D_i^a(h_1))Y_i(h_1; 0). \quad (17)$$

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 $Y_{i,2}^a$.

⁴¹This 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 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.

⁴²The 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.

⁴³In 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.

⁴⁴Since 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.

We are interested in the effects of the policy on the present value of earnings, so let $\tilde{Y}_i^a \equiv Y_{i,1}^a + \rho Y_{i,2}^a$, 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 $E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a]$. 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 could 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.^{45,46} 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.⁴⁷ All in all, the effect of the policy change on the present value of earnings is given by $\Delta_{a,a'} = E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a \mid D_i^{a'}(h_1) = 1, D_i^a(h_1) = 0]$, the dynamic policy relevant treatment parameter (*DP RTE*). 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 decomposed as:

$$\begin{aligned}
\Delta_{a,a'} = & \underbrace{E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a \mid D_i^{a'}(1) = D_i^a(0) = 1, \mathbb{A}_i]}_{\text{Compliers, Always-Takers } (W_{CO,AT})} \times \underbrace{\Pr[D_i^{a'}(1) = D_i^a(0) = 1 \mid \mathbb{A}_i]}_{\text{Probability of Always-Takers}} \\
& + \underbrace{E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a \mid D_i^{a'}(1) = 1, D_i^a(0) = 0, \mathbb{A}_i]}_{\text{Compliers, Compliers } (W_{CO,CO})} \times \underbrace{\Pr[D_i^{a'}(1) = 1, D_i^a(0) = 0 \mid \mathbb{A}_i]}_{\text{Probability of Compliers}} \\
& + \underbrace{E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a \mid D_i^{a'}(1) = 0, D_i^a(0) = 0, \mathbb{A}_i]}_{\text{Compliers, Never-Takers } (W_{CO,NT})} \times \underbrace{\Pr[D_i^{a'}(1) = 0, D_i^a(0) = 0 \mid \mathbb{A}_i]}_{\text{Probability of Never-Takers}} \\
& + \underbrace{E[\tilde{Y}_i^{a'} - \tilde{Y}_i^a \mid D_i^{a'}(1) = 0, D_i^a(0) = 1, \mathbb{A}_i]}_{\text{Compliers, Defiers } (W_{CO,DF})} \times \underbrace{\Pr[D_i^{a'}(1) = 0, D_i^a(0) = 1 \mid \mathbb{A}_i]}_{\text{Probability of Defiers}} \tag{18}
\end{aligned}$$

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

⁴⁵We note that 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 a static world, where only contemporary compliers are the relevant group (Heckman and Vytlacil, 2001; Mogstad and Torgovitsky, 2018).

⁴⁶In 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.

⁴⁷In the notation of Blanco et al. (2019), these are our “principal strata”, as we could compare outcomes within groups with different treatment status to recover treatment effects. In contrast to the principal stratification literature, we follow a point-identification approach that exploits exclusion restrictions and a factor structure.

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.

6.3 Results

Table 7: 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%	Policy: +50%
	(1)	(2)
Compliers (CO)	0.28%	1.45%
Weights by Type		
Compliers-Always Takers (<i>CO, AT</i>)	0.099	0.105
Compliers-Compliers (<i>CO, CO</i>)	0.304	0.299
Compliers-Never Takers (<i>CO, NT</i>)	0.495	0.485
Compliers-Defiers (<i>CO, DF</i>)	0.103	0.111

Panel B. Dynamic Policy Relevant Treatment Effect by Response Type

Dynamic Policy Relevant Treatment Effects		
	Policy: +10%	Policy: +50%
	(1)	(2)
Compliers	3.85%	4.12%
	[2.96,4.74]	[3.74,4.50]
Compliers-Always Takers (<i>CO, AT</i>)	3.32%	2.93%
	[0.63,6.00]	[1.79,4.07]
Compliers-Compliers (<i>CO, CO</i>)	4.57%	4.67%
	[2.99,6.15]	[3.98,5.36]
Compliers-Never Takers (<i>CO, NT</i>)	3.66%	4.27%
	[2.35,4.96]	[3.72,4.82]
Compliers-Defiers (<i>CO, DF</i>)	3.18%	3.11%
	[0.46,5.89]	[2.00,4.22]

Notes: Table 7 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 $\tilde{Y}_i^{aj} \equiv Y_{i1}^{aj} + \rho Y_{i2}^{aj}$ be the present value of earnings for a policy regime a_j . Weights are as defined in Section 5.

The first panel of Table 7 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 0.3% 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 1.5% of young workers to change their training decision.⁴⁸ 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

⁴⁸While 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.

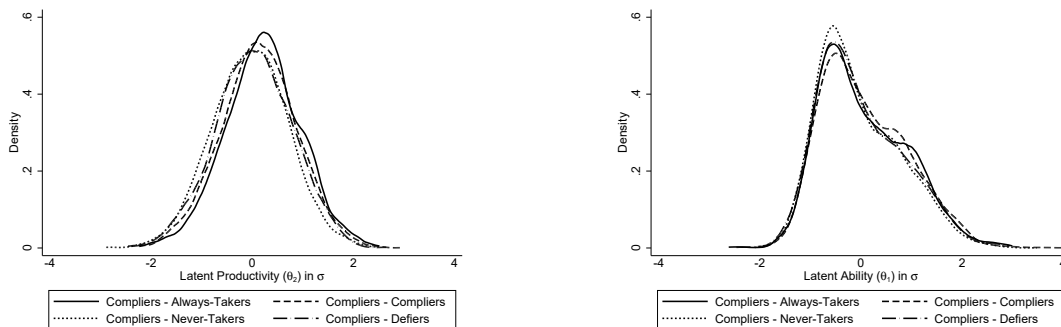
period in response to the policy. Meanwhile, the complier-complier group, who take-up training in both time periods due to the policy change, accounts for 30% of all affected workers. The weights assigned to dynamic response types are largely constant across the two interventions.

In the second panel of Table 7, 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 3.8%, reaching 4.1% for a 50 percent increase in course-hour availability. While the impacts are captured by the *DP RTE* parameter, the effect sizes fit largely in line with the dynamic average treatment effect (*DATE*) presented in Table 6.

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 4.6% range. We find similar effects for workers who are induced to participate in *only* one training course, reaching 3.6-4.2% for compliers-never-takers. The complier-defier group, which captures workers who move up their training choice due to the policy change, would face returns in the 3.1% range, which are largely similar to the returns for the complier-always taker group.

In Figure 9, we explore whether differences in the density of the unobserved factors across response type groups may account for the heterogeneous earnings impacts documented above. The latent productivity density of compliers-always-takers dominates that of the other response types, surpassing that of compliers-compliers by 0.11 standard deviations, on average. Moreover, there are larger differences with the other response types, as the average latent productivity of compliers-always takers exceeds that of compliers-never takers by 0.32 standard deviations. In the second panel, we present similar patterns along the latent ability (θ_1) margin, yet the differences are smaller in magnitude, as compliers-always takers surpass compliers-never takers by just 0.11 σ in this dimension.

Figure 9: Density of Latent Factors by Dynamic Response Types



(a) Response Types by Latent Productivity (θ_2)

(b) Response Types by Latent Ability (θ_1)

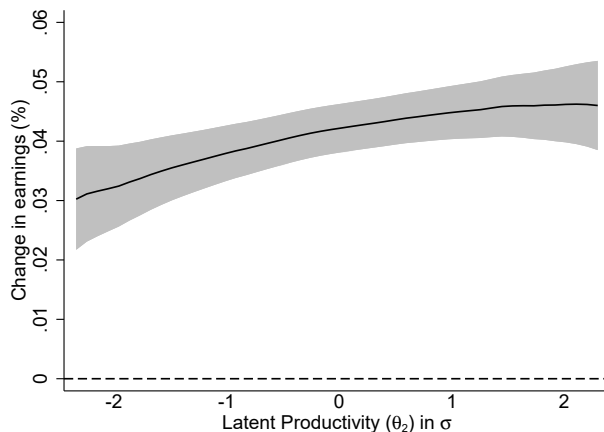
Note: Figure 9 shows the estimated density of latent productivity (θ_2) and latent ability (θ_1) across dynamic response types. The simulated policy reflects a 50% expansion in program course-hours availability in both periods. The density of the latent factor distribution across response types is similar across different program expansion levels.

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 θ_2 (Figure 10).⁴⁹ We find largely

⁴⁹The 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, in 2015 the Chilean government 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

homogeneous effects across the ability distribution, in the range of 3.8%.⁵⁰

Figure 10: Heterogeneous Dynamic Policy Relevant Treatment Effects



Note: Figure 10 shows the estimated impact of a 50% expansion in course availability on the present-value of earnings across levels of latent productivity (θ_2), captured by the dynamic policy-relevant treatment effect parameter.

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 along two dimensions of unobserved heterogeneity, remarking differences in sorting-into-training on both the timing and prevalence of participation in the early career. 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 factor 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 (substitutability). We find dynamic substitutability of first-year training: early investments decrease the economic returns to later investments. 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 formally test for potential mechanisms, we consider our results a first step towards understanding the complex dynamic of the returns to training.

cost than a small expansion.

⁵⁰Along the latent ability component, the DPRTE is higher for lower skilled workers, resembling the DATE patterns presented in Figure 7.

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.
- Ashworth, J., V. J. Hotz, A. Maurel, and T. Ransom (2017). Changes across cohorts in wage returns to schooling and early work experiences. Technical report, National Bureau of Economic Research.
- Attanasio, O., R. Blundell, G. Conti, and G. Mason (2020). Inequality in socio-emotional skills: A cross-cohort comparison. *Journal of Public Economics*, 104171.
- Attanasio, O., S. Cattan, E. Fitzsimons, C. Meghir, and M. Rubio-Codina (2020). Estimating the

- production function for human capital: results from a randomized controlled trial in Colombia. *American Economic Review* 110(1), 48–85.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). The skill content of recent technological change: An empirical exploration. *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.
- Blanco, G., X. Chen, C. A. Flores, and A. Flores-Lagunes (2019). Bounds on average and quantile treatment effects on duration outcomes under censoring, selection, and noncompliance. *Journal of Business & Economic Statistics*, 1–38.
- Booth, A. L. (1993). Private sector training and graduate earnings. *The Review of Economics and Statistics* 75, 164–170.
- Callaway, B. and P. H. Sant’Anna (2020). Difference-in-differences with multiple time periods. *Journal of Econometrics*.
- 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). 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.
- 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, pp. 697–812. Elsevier.
- De Chaisemartin, C. and X. d’Haultfoeulle (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110(9), 2964–96.
- 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).
- Fruehwirth, J. C., S. Navarro, and Y. Takahashi (2016). How the timing of grade retention affects outcomes: Identification and estimation of time-varying treatment effects. *Journal of Labor Economics* 34(4), 979–1021.
- Goodman-Bacon, A. (2018). Difference-in-Differences with Variation in Treatment Timing. Working Paper.
- Guvenen, F., B. Kuruscu, S. Tanaka, and D. Wiczer (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics* 12(1), 210–44.
- 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. and S. Urzua (2010). Comparing iv with structural models: What simple iv can and cannot identify. *Journal of Econometrics* 156(1), 27–37.
- 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. Mosso (2014). The economics of human development and social mobility. *Annu. Rev. Econ.* 6(1), 689–733.
- 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 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 (2019). Reducing inequality through dynamic complementarity: Evidence from head start and public school spending. *American Economic Journal: Economic Policy* 11(4), 310–49.
- 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.
- 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.
- Lise, J. and F. Postel-Vinay (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review* 110(8), 2328–76.
- Lynch, L. M. (1992). Private-sector training and the earnings of young workers. *American Economic Review* 82, 299–312.
- Matzkin, R. (1994). *Restrictions of Economic Theory in Nonparametric Methods*. New York: North-Holland.
- 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 parameters. *Econometrica* 86(5), 1589–1619.
- Mogstad, M. and A. Torgovitsky (2018). Identification and extrapolation of causal effects with instrumental variables. *Annual Review of Economics* 10, 577–613.
- 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.
- 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.
- 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.
- Spitz-Oener, A. (2006, apr). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics* 24(2), 235–270.
- Urzua, S. (2008). Racial labor market gaps the role of abilities and schooling choices. *Journal of Human Resources* 43(4), 919–971.
- Veum, J. (1997). Training and job mobility among young workers in the united states. *Journal of Population Economics* 10, 219–233.

- SUPPLEMENTARY APPENDICES -

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.

We follow the structure of the model to define the average treatment effect parameter. The impact of training for an individual i at period t for a given history h equals $Y_{it}(h_t; 1) - Y_{it}(h_t; 0)$.^{51,52} 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:

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

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:

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

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

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

If the data generating process follows the dynamic model introduced in Section 2, 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:

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

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

⁵¹We 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} .

⁵²For notational simplicity, let $X^Y \beta^Y(h_t; j) = \mu^Y(h_t; j)$.

are given by:

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

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

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

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

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:

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

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} .⁵³ 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 $t \in \mathcal{T}$.

⁵³When $T = 2$, 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:*

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

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). □

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:

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

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).⁵⁴ 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.

⁵⁴We present the test for the latent ability factor θ_1 . Results are equivalent for θ_2 .

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	0.025	0.040	0.010
p-value of test	0.000		

(b) Assumption 2: $\lambda^Y(h_t; 1) - \lambda^Y(h_t; 0)$			
	$\lambda^Y(1) - \lambda^Y(0)$	$\lambda^Y(0, 1) - \lambda^Y(0, 0)$	$\lambda^Y(1, 1) - \lambda^Y(1, 0)$
Coefficient	-0.010	0.004	0.013
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.

B Identification Strategy

This section presents the identification of the measurement system presented in Section 4. The identification of the distribution of the Latent Factors follows the formal arguments presented in [Carneiro et al. \(2003\)](#), [Hansen et al. \(2004\)](#), [Heckman et al. \(2006\)](#). In the measurement system presented in equations Section 4, the covariance between all test scores is observed and relied on as part of the identification strategy. Throughout this section, we keep the conditioning on \mathbf{X} implicit. Let T_k denote observed test score measures ($k = 1, 2, 3$). Using the covariances between these test scores, we can compute:

$$\begin{aligned} Cov(T_1, T_2) &= \lambda_1^T \lambda_2^T \sigma_{\theta,1}^2 \\ Cov(T_2, T_3) &= \lambda_2^T \lambda_3^T \sigma_{\theta,1}^2 \\ Cov(T_1, T_3) &= \lambda_1^T \lambda_3^T \sigma_{\theta,1}^2 \end{aligned}$$

where $\sigma_{\theta,1}^2$ represents the variance of the latent ability factor. Normalizing the loading associated with the PSU math test score ($\lambda_1^T = 1$), yields a system with three equations and three unknowns. We can identify the remaining three unknown parameters λ_2^T , λ_3^T , and $\sigma_{\theta,1}^2$.

For the initial wage system, we identify the latent ability factor loadings leveraging the covariance of test scores and wages as follows:

$$\begin{aligned} Cov(T_1, W_1) &= \lambda_1^T \lambda_1^W \sigma_{\theta,1}^2 \\ Cov(T_1, W_2) &= \lambda_1^T \lambda_2^W \sigma_{\theta,1}^2 \end{aligned}$$

Since λ_1^T and $\sigma_{\theta,1}^2$ are already identified, the remaining loadings λ_1^W and λ_2^W are also identified from each equation presented above.

To identify the variance of the latent productivity factor, we take advantage of the first two initial wages along with the first training probit (equation (5)), as in [Urzua \(2008\)](#). The covariances across these measures are given by:

$$\begin{aligned} Cov(W_1, W_2) &= \lambda_1^W \lambda_2^W \sigma_{\theta,1}^2 + \eta_1^W \eta_2^W \sigma_{\theta,2}^2 \\ Cov(W_1, I) &= \lambda_1^W \lambda^I \sigma_{\theta,1}^2 + \eta_1^W \eta^I \sigma_{\theta,2}^2 \\ Cov(W_2, I) &= \lambda_2^W \lambda^I \sigma_{\theta,1}^2 + \eta_2^W \eta^I \sigma_{\theta,2}^2 \end{aligned}$$

where $\sigma_{\theta,2}^2$ represents the variance of the latent productivity factor. As the λ_j^W components are already identified, the system above includes three equations and four unknowns. By normalizing the loading associated with the first initial wage ($\eta_1^W = 1$), the remaining loadings (η_1^W and η^I) and the variance of the latent productivity factor ($\sigma_{\theta,2}^2$) are identified as well. Having secured the identification of all the loadings and the variance of each latent component, we can apply the following transformation to the test score measurement system:

$$\frac{T_k}{\lambda_k^T} = \theta_1 + \frac{\varepsilon_{e_k}}{\lambda_1^T} \tag{B.1}$$

We can then apply [Kotlarski \(1967\)](#)'s theorem to equation (B.1) to identify:

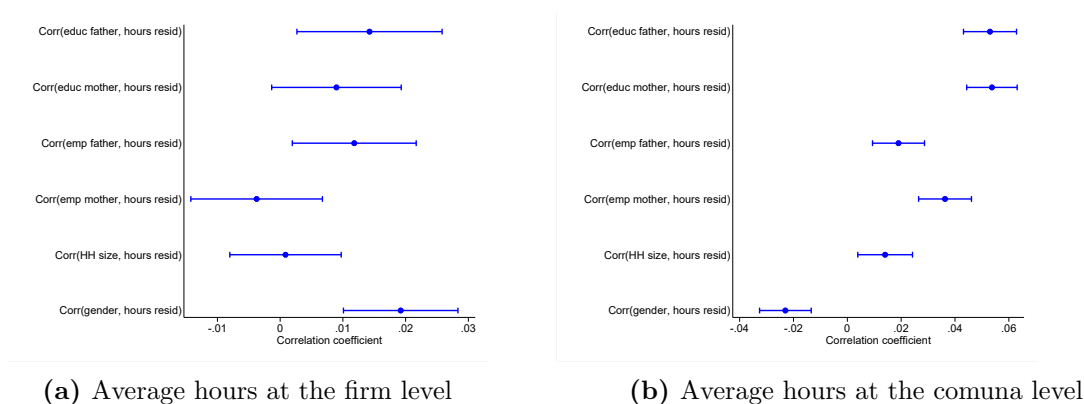
$$f_{\theta_1}(\cdot), f_{\varepsilon_{T_k}}(\cdot) \tag{B.2}$$

Applying the same argument to equation (10) allows us to identify $f_{\theta_2}(\cdot), f_{\varepsilon_{W_j}}(\cdot)$

C Evaluating the role of Exclusion Restrictions

C.1 Instruments Diagnosis

Figure C.1: Instruments Diagnosis: Correlations between Instruments and Parents Education



Note: This figure shows estimated correlation coefficients of our instrument with parents education and of these variables with average wages. The left panel shows estimated correlations of hours at the firm while the right panel at the comuna level.

C.2 Returns to Job Training: No Instrument Model

Table C.1: Static Returns to Training


Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	-0.06 [-0.09,-0.02]	3.66 [3.58,3.75]	0.50 [0.35,0.66]
<i>TT</i> (percentage points)	-0.09 [-0.16,-0.01]	3.87 [3.64,4.11]	0.59 [0.47,0.70]

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.2: Dynamic Returns to Training

	DATE	DTT	DTUT
Direct effect (short-term)	0.30 (0.01) [16%]	0.29 (0.02) [16%]	0.31 (0.01) [16%]
Direct effect (medium-term)	1.60 (0.02) [85%]	1.65 (0.04) [89%]	1.59 (0.02) [85%]
Continuation value	-0.03 (0.02) [-2%]	-0.09 (0.04) [-5%]	-0.01 (0.02) [-1%]
Total	1.88 (0.02)	1.86 (0.05)	1.88 (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.

Figure C.2: Goodness of Fit: Training Decisions

Note: In Figure C.2, 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.

D The Gibbs Sampler Algorithm for a Roy Model

We describe the our MCMC algorithm. For sake of clarity, we do this in the context of a simplified version of the Roy model: a two-sector static one-factor model. Its extension to the setting used in the main text is direct. Let θ be the latent factor and consider the following model:

$$\begin{aligned} I &= Z\gamma + V, \\ Y_1 &= X_0\beta_{1,0} + X_1\beta_{1,1} + \varepsilon_1, \\ Y_0 &= X_0\beta_{0,0} + X_1\beta_{0,1} + \varepsilon_0. \\ D &= 1[I > 0] \end{aligned}$$

Assume a factor structure of the form:

$$\begin{aligned} V &= \theta\alpha_V + U_V \\ \varepsilon_1 &= \theta\alpha_1 + U_1 \\ \varepsilon_0 &= \theta\alpha_0 + U_0 \end{aligned}$$

where $U_1 \perp U_0 \perp U_V$, and $\theta \perp (U_0, U_1, U_V)$, and

$$\begin{aligned} U_1 &\sim N(0, \sigma_{U_1}^2), \\ U_0 &\sim N(0, \sigma_{U_0}^2), \\ U_V &\sim N(0, \sigma_{U_V}^2). \end{aligned}$$

As a measurement system, consider the test score equation:

$$T = Q\kappa + \theta + U_{T_1}$$

where $U_{T_1} \sim N(0, \sigma_{U_{T_1}}^2)$.

Thus, the likelihood function can be written as:

$$\begin{aligned} f(Y_j, T_j, D_j; \Theta) &= \int f(Y_j, T_j, D_j | \theta_j) dF(\theta_j) \\ &= \int f(Y_j, D_j | \theta_j) f(T_j | \theta_j) dF(\theta_j). \end{aligned}$$

Under our assumptions:

$$\begin{aligned} \Gamma(Y, T, D; \Theta) &= \prod_{j=1}^N \int \left[\frac{1}{\sqrt{2\pi}\sigma_T} \exp\left(-\frac{1}{2\sigma_T^2}(T_j - Q_j\kappa - \theta_j)^2\right) \right] \\ &\quad \left[\frac{1}{\sqrt{2\pi}\sigma_{u_1}} \exp\left(-\frac{1}{2\sigma_{u_1}^2}(Y_{1,j} - X_j\beta_1 - \alpha_1\theta_j)^2\right) \Phi(-Z_j\gamma - \alpha_V\theta_j) \right]^{D_j} \\ &\quad \left[\frac{1}{\sqrt{2\pi}\sigma_{u_0}} \exp\left(-\frac{1}{2\sigma_{u_0}^2}(Y_{0,j} - X_j\beta_0 - \alpha_0\theta_j)^2\right) (1 - \Phi(-Z_j\gamma - \alpha_V\theta_j)) \right]^{1-D_j} dF(\theta_j). \end{aligned}$$

The block structure associated with the likelihood function becomes:

$$\begin{aligned}
f(\alpha, \beta, \tau, \gamma, \kappa, \theta/Y, T, D) &\propto f(Y, T, D|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(\alpha, \beta, \tau, \gamma, \kappa, \theta) \\
&\propto f(Y, D|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(T|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(\alpha, \beta, \tau, \gamma, \kappa, \theta) \\
&\propto f(Y|D, \alpha, \beta, \tau, \gamma, \kappa, \theta)f(D|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(T|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(\alpha, \beta, \tau, \gamma, \kappa, \theta) \\
&\propto f(Y|D, \alpha, \beta, \tau, \gamma, \kappa, \theta)f(D|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(T|\alpha, \beta, \tau, \gamma, \kappa, \theta)f(\alpha, \beta, \tau, \gamma, \kappa, \theta) \\
&= \tau_{u1}^{n_{(D_i=1)}/2} \exp\left(-\tau_{u1} \sum_{(D_i=1)} (Y_i^1 - X_i\beta_1 - \alpha_1\theta_i)^2\right) \\
&\quad \times \tau_{u0}^{n_{(D_i=0)}/2} \exp\left(-\tau_{u0} \sum_{(D_i=0)} (Y_i^0 - X_i\beta_0 - \alpha_0\theta_i)^2\right) \\
&\quad \times \prod_{i=1}^n [\Phi(-Z_i\gamma - \alpha_v\theta_i)]^{D_i} [(1 - \Phi(-Z_i\gamma - \alpha_v\theta_i))]^{1-D_i} \\
&\quad \times \tau_{ut}^{n/2} \exp\left(-\tau_{ut}(\sum (T_i - Q_i\kappa - \alpha_t\theta_i)^2)\right) \\
&\quad \times \tau_f^{n/2} \exp(-\tau_f(\sum \theta_i^2)) \exp(-\tau_{u1}) \exp(-\tau_{u0}) \exp(-\tau_{ut}) \\
&\quad \times \left(\frac{1}{10}\right)^{1/2} \exp\left\{-\frac{1}{10}(\alpha_0)^2\right\} \left(\frac{1}{10}\right)^{1/2} \exp\left\{-\frac{1}{10}(\alpha_1)^2\right\} \\
&\quad \left(\frac{1}{10}\right)^{1/2} \exp\left\{-\frac{1}{10}(\alpha_v)^2\right\},
\end{aligned}$$

where we explicitly impose a set of prior distributions. Using the block structure we can derive the formulae for the conditional posteriors:

1. Outcome equations:

$$f(\beta_i/\alpha_i, \tau_i, \theta, Y, D) \propto \exp\left\{-\tau_i \left(\sum_{j:D=i} (Y_j - X_j\beta_i - \alpha_i\theta_j)^2\right)\right\} \text{ for } i = 0, 1$$

so

$$\beta_i/. \sim N\left(\frac{\sum_{j:D=i} x_j(y_j - \alpha_i\theta_j)}{\sum_{j:D=i} x_j x'_j}, \left(\tau_i \sum_{j:D=i} x_j x'_j\right)^{-1}\right), \quad (\text{D.1})$$

$$f(\alpha_i/\beta_i, \tau_i, f, Y_i) \propto \exp\left\{-\frac{1}{2}\tau_i \sum_{j:D=i} (y_j - x_j\beta_i - \alpha_i\theta_j)^2 - \frac{1}{2}\frac{1}{10}\alpha_i^2\right\} \text{ for } i = 0, 1.$$

Let $\tilde{Y} = Y - X\beta$, thus

$$\exp\left\{-\frac{1}{2}\tau_{u_i} \sum_{j:D=i} (\tilde{y}_j - \alpha_i\theta_j)^2 - \frac{1}{2}\frac{1}{10}\alpha_i^2\right\}$$

Thus, we can demonstrate:

$$\alpha_i/\beta_i, \tau_i, \theta, Y, D \sim N(\hat{\alpha}_i, \bar{\Sigma}_i) \quad (\text{D.2})$$

where $\bar{\Sigma}_i = \left(\tau_i \sum_{j:D=i} \theta_j^2 + \frac{1}{10} \right)^{-1}$ and $\hat{\alpha}_i = \bar{\Sigma}_i \left(\tau_i \sum_{j:D=i} \theta_j \tilde{y}_j \right)$.

2. Measurement system:

$$f(\kappa/\tau_T, \theta, Q, T) \propto \exp \left(-\tau_T \sum_{j=1}^N (T_j - Q_j \kappa - \theta_j)^2 \right)$$

then

$$\kappa/. \sim N((Q'Q)^{-1} (Q'(T - \theta)), \tau_{ut}^{-1} (Q'Q)^{-1}) \quad (\text{D.3})$$

3. Decision model:

We need a completion for D . Let D^* be the latent variable. Then,

$$\begin{aligned} f(\gamma, \alpha_V, D^*/D) &\propto f(D/\alpha_V, \gamma, D^*) f(D^*/\alpha_V, \gamma) f(\gamma) f(\alpha_V) \\ &= \left(\frac{1}{10} \right)^{1/2} \exp \left\{ -\frac{1}{2} \frac{1}{10} (\alpha_V)^2 \right\} \prod_{j=1}^N [1(D_j^* > 0) \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j)]^{D_j} \\ &\quad [1(D_j^* < 0) \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j)]^{1-D_j} \end{aligned}$$

Now,

$$\begin{aligned} f(\gamma/\alpha_V, \theta, D^*, D) &\propto \prod_{j=1}^N \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j) \\ &= \exp \left\{ \frac{1}{2} \sum_{j=1}^N (D_j^* - Z_j \gamma - \alpha_V \theta_j)^2 \right\} \end{aligned}$$

Consequently,

$$\gamma/. \sim N((Z'Z)^{-1} (Z'(D^* - \theta \alpha_V)), (Z'Z)^{-1}) \quad (\text{D.4})$$

On the other hand,

$$\begin{aligned} f(\alpha_V/\gamma, D^*, D) &\propto \exp \left\{ -\frac{1}{10} (\alpha_V)^2 \right\} \prod_{j=1}^N \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j) \\ &= \exp \left\{ \frac{1}{2} \sum_{j=1}^N (D_j^* - Z_j \gamma - \alpha_V \theta_j)^2 - \frac{1}{2} \frac{1}{10} (\alpha_V)^2 \right\} \end{aligned}$$

Using the same logic as in (D.2) we obtain:

$$\alpha_V/. \sim N(\hat{\alpha}_V, \bar{\Sigma}_V) \quad (\text{D.5})$$

where $\bar{\Sigma}_V = \left(\theta^2 + \frac{1}{10} \right)^{-1}$ and $\hat{\alpha}_V = \bar{\Sigma}_V (\theta (D^* - Z\gamma))$.

Finally,

$$f(D^*/\alpha_V, \theta, \gamma, D) \propto [1(D_j^* > 0) \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j)]^{D_j} [1(D_j^* < 0) \phi(D_j^* - Z_j \gamma - \alpha_V \theta_j)]^{1-D_j}$$

Therefore, we sample D^* from:

$$D_j^* = \begin{cases} TN_{[0,\infty)}(Z_j\gamma + \alpha_v\theta_j, 1) & \text{if } D_j = 1 \\ TN_{(-\infty,0]}(Z_j\gamma + \alpha_v\theta_j, 1) & \text{if } D_j = 0 \end{cases} \quad (\text{D.6})$$

4. Precisions:

$$\tau_{i/} \sim G \left(\frac{\sum_{j:D=i} 1}{2} + 2, \left(\frac{1}{2} \sum_{j:D=i} (y_j - x_j\beta_i - \alpha_i\theta_j)^2 \right) + 1 \right) \quad \text{for } i = 0, 1 \quad (\text{D.7})$$

$$\tau_{T/} \sim G \left(\frac{N}{2} + 2, \left(\frac{1}{2} \sum_{j=1}^N (T_j - Q_j\kappa - \theta_j)^2 \right) + 1 \right) \quad (\text{D.8})$$

$$\tau_{\theta/} \sim G \left(\frac{N}{2} + 2, \sum_{j=1}^N \theta_j^2 + 1 \right) \quad (\text{D.9})$$

Lastly, the posterior for θ becomes:

$$\begin{aligned} \theta_{j/}, D_j = 1 &\propto \tau_i^{1/2} \exp \left(-\tau_i \frac{1}{2} (Y_j - X_j\beta_i - \alpha_i\theta_j)^2 \right) \\ &\tau_T^{1/2} \exp \left(-\tau_{ut} \frac{1}{2} (T_j - Q_j\kappa - \theta_j)^2 \right) \exp \left(\frac{1}{2} (D_j^* - Z_j\gamma - \alpha_v\theta_j)^2 \right) \\ &\tau_f^{1/2} \exp(-\tau_f\theta_j^2) \\ &\propto \exp(\tau_i(Y_j - X_j\beta_i - \alpha_i\theta_j)^2) \exp(\tau_T(T_j - Q_j\kappa - \theta_j)^2) \\ &\quad \exp \left((D_j^* - Z_j\gamma - \alpha_v\theta_j)^2 \right) \exp(\tau_f\theta_j^2) \\ &= \exp \left(\tau_i(\theta_j^* - \alpha_i\theta_j)^2 + \tau_{ut}(\theta_j^{**} - \theta_j)^2 + (\theta_j^{***} - \alpha_v\theta_j)^2 + \tau_f\theta_j^2 \right), \end{aligned}$$

where $\theta_j^* = Y_j - X_j\beta_i$, $\theta_j^{**} = T_j - Q_j\kappa$, $\theta_j^{***} = D_j^* - Z_j\gamma$. Notice that:

$$\begin{aligned} \alpha_i^2 \tau_i \left(\frac{1}{\alpha_i} \theta_j^* - \theta_j \right)^2 + \alpha_t^2 \tau_T (\theta_j^{**} - \theta_j)^2 &= (\alpha_i^2 \tau_{u_i} + \alpha_t^2 \tau_{ut}) \left(f_j - \frac{\alpha_i \tau_{u_i} \theta_j^* + \alpha_t \tau_{ut} \theta_j^{**}}{(\alpha_i^2 \tau_{u_i} + \alpha_t^2 \tau_{ut})} \right)^2 \\ &\quad + \frac{(\alpha_i^2 \tau_{u_i})(\alpha_t^2 \tau_{ut})}{(\alpha_i^2 \tau_{u_i} + \alpha_t^2 \tau_{ut})} \left(\frac{1}{\alpha_i} \theta_j^* - \frac{1}{\alpha_t} \theta_j^{**} \right)^2, \end{aligned}$$

and

$$\alpha_v^2 \left(\frac{1}{\alpha_v} \theta_j^{***} - \theta_j \right)^2 + \tau_f \theta_j^2 = (\alpha_v^2 + \tau_f) \left(\theta_j - \frac{\alpha_v \theta_j^{***}}{(\alpha_v^2 + \tau_f)} \right)^2 - \frac{\alpha_v^2 \tau_f}{(\alpha_v^2 + \tau_f)} \left(-\frac{1}{\alpha_v} \theta_j^{***} \right)^2.$$

Thus,

$$\begin{aligned} \exp\left(\tau_i(\theta_j^* - \alpha_i\theta_j)^2 + \tau_T(\theta_j^{**} - \theta_j)^2 + (\theta_j^{***} - \alpha_V\theta_j)^2 + \tau_f\theta_j^2\right) \propto \\ \exp\left((\alpha_i^2\tau_i + \tau_T)\left(\theta_j - \frac{\alpha_i\tau_i\theta_j^* + \tau_T\theta_j^{**}}{(\alpha_i^2\tau_i + \tau_T)}\right)^2\right. \\ \left. + (\alpha_V^2 + \tau_f)\left(\theta_j - \frac{\alpha_V\theta_j^{***}}{(\alpha_V^2 + \tau_f)}\right)^2\right) \end{aligned}$$

And since:

$$\begin{aligned} \exp\left(\tau_i(\theta_j^* - \alpha_i\theta_j)^2 + \tau_T(\theta_j^{**} - \theta_j)^2 + (\theta_j^{***} - \alpha_V\theta_j)^2 + \tau_f\theta_j^2\right) \propto \\ \exp\left(-\frac{1}{2}(\alpha_i^2\tau_i + \tau_T + \alpha_V + \tau_f)\left(\theta_j - \frac{\alpha_i\tau_i\theta_j^* + \tau_T\theta_j^{**} + \alpha_V\theta_j^{***}}{(\alpha_i^2\tau_i + \tau_T + \alpha_V + \tau_f)}\right)^2\right), \end{aligned}$$

we obtain:

$$\theta_j/., D_j = i \sim N\left(\frac{\alpha_i\tau_{u_i}\theta_j^* + \alpha_t\tau_{ut}\theta_j^{**} + \alpha_v\theta_j^{***}}{(\alpha_i^2\tau_{u_i} + \alpha_t^2\tau_{ut} + \alpha_v^2 + \tau_\theta)^{-1}}, (\alpha_i^2\tau_{u_i} + \alpha_t^2\tau_{ut} + \alpha_v^2 + \tau_\theta)^{-1}\right). \quad (\text{D.10})$$

The Gibbs sampling procedure becomes:

1. Choose initial values for the parameters, and an arbitrary first draw for the factor. For example, $\theta^{(m)} \sim N(0, 1)$

For $m = 1, M$

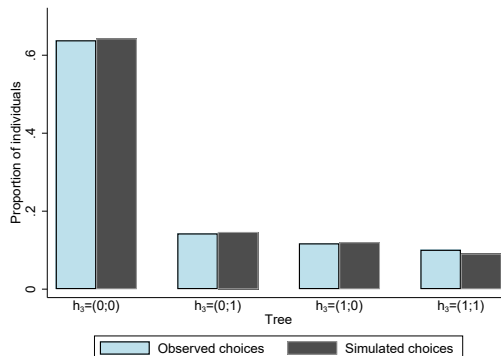
1. Sample $D_j^{*(m)}$ for $j = 1, \dots, N$ according to (D.6)
2. Sample $\theta_j^{(m)}$ for $j = 1, \dots, N$ according to (D.10)
3. Sample $\beta_i^{(m)}$ ($i = 1, 2$) according to (D.1)
4. Sample $\alpha_i^{(m)}$ ($i = 1, 2$) according to (D.2)
5. Sample $\kappa^{(m)}$ according to (D.3)
6. Sample $\gamma^{(m)}$ according to (D.4)
7. Sample $\alpha_V^{(m)}$ according to (D.5)
8. Sample $\tau_i^{(m)}$ ($i = 1, 2$) according to (D.7)
9. Sample $\tau_T^{(m)}$ according to (D.8)
10. Sample $\tau_f^{(m)}$ according to (D.9)

Iterate over m until converge.

E Estimated Model Parameters

E.1 Goodness of Fit

Figure E.1: Goodness of Fit: Training Decisions



Note: In Figure E.1, 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.

Table E.1: Goodness of Fit: Labor Market Outcomes by Training History

Moment	Actual	Model	H_0 : Actual=Model p-value
A. Period $t = 1$			
Unconditional: $E[Y_1]$	6.162	6.166 [6.161,6.171]	0.244
Conditional on training decision: $E[Y(h_1; j) j = 1]$	6.304	6.245 [6.234,6.256]	0.000
$E[Y(h_1; j) j = 0]$	6.122	6.144 [6.138,6.150]	0.000
B. Period $t = 2$			
Unconditional: $E[Y_2]$	6.273	6.277 [6.271,6.282]	0.351
Conditional on training decision: $E[Y(h_2; j) h_2 = \{1\}, j = 1]$	6.550	6.441 [6.423,6.458]	0.000
$E[Y(h_2; j) h_2 = \{1\}, j = 0]$	6.292	6.291 [6.275,6.306]	0.893
$E[Y(h_2; j) h_2 = \{0\}, j = 1]$	6.405	6.374 [6.360,6.389]	0.004
$E[Y(h_2; j) h_2 = \{0\}, j = 0]$	6.196	6.228 [6.221,6.235]	0.000

Notes: Table E.1 shows the means of log earnings by year and training choice from the observed data and the simulated sample. Following the notation used in the text, $Y(h_t, j)$ represents (log) earnings in period t given training history h_t and training decision at that period j . Y_t^a represents unconditional earnings reported at t . Thus, $Y_t = Y(h_t, 1) \times D(h_t) + Y(h_t, 0) \times (1 - D(h_t))$. In brackets, we show a 95% confidence interval on the mean of simulated earnings by training decision. We present the p-value from a t -test of the equality of means.

E.2 Estimated Parameters

Table E.2: Measurement system estimates

	Math PSU	Language PSU	HS GPA	Salary (I)	Salary (II)
Constant	2.204 (0.117)	1.774 (0.122)	3.945 (0.135)	-3.946 (0.062)	-2.954 (0.062)
Male	0.134 (0.009)	-0.083 (0.009)	-0.383 (0.010)	0.119 (0.010)	0.018 (0.010)
Age	-0.156 (0.007)	-0.127 (0.007)	-0.208 (0.007)		
Household Size	-0.021 (0.002)	-0.037 (0.003)	0.004 (0.003)		
Mother's Education	0.038 (0.002)	0.044 (0.002)	0.004 (0.002)		
Father's Education	0.027 (0.001)	0.029 (0.001)	-0.003 (0.001)		
Father's Employment	-0.097 (0.011)	-0.104 (0.012)	-0.066 (0.012)		
Mother's Employment	-0.011 (0.009)	0.025 (0.010)	-0.110 (0.011)		
Productivity Factor	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	1.000 (0.000)	0.719 (0.006)
Ability Factor	1.000 (0.000)	0.852 (0.007)	0.565 (0.007)	0.222 (0.007)	0.271 (0.007)
Precision	6.867 (0.227)	2.849 (0.035)	1.453 (0.012)	3.850 (0.032)	2.011 (0.016)
Age at Entry				0.167 (0.003)	0.137 (0.003)
Sample Size	37,089				

Note: The table displays the estimation results from the measurement system of test scores (equation 9). We obtain these estimates by simulating 500 values of parameters using our estimated posterior. The 'Sample Size' row denotes the number of individuals included in our data and used to estimate the model. 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 E.3: Structural Model: training probits

	I	$I(0)$	$I(1)$
Constant	-1.323 (0.094)	-1.546 (0.106)	-1.342 (0.174)
Gender	-0.037 (0.015)	-0.037 (0.018)	-0.082 (0.030)
Firm average training hours	0.605 (0.010)	0.722 (0.012)	0.351 (0.017)
Comuna average training hours	0.164 (0.085)	0.104 (0.091)	-0.071 (0.164)
Comuna Wages	0.024 (0.007)	0.011 (0.009)	0.004 (0.015)
Productivity Factor	0.090 (0.012)	0.141 (0.013)	0.222 (0.020)
Ability Factor	0.030 (0.011)	0.029 (0.012)	0.087 (0.020)
Precision	1.000 (0.000)	1.000 (0.000)	1.000 (0.000)
Sample Size	37,089		

Note: We show the estimated parameters of the training probits (equation 5). We obtain these estimates by simulating 500 values of parameters using our estimated posterior. The ‘Sample Size’ row denotes the number of individuals included in our data and used to estimate the model. 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 E.4: Structural model: earnings equations

	$Y(0)$	$Y(1)$	$Y(0, 0)$	$Y(0, 1)$	$Y(1, 0)$	$Y(1, 1)$
Constant	4.788 (0.156)	3.632 (0.312)	4.668 (0.220)	3.868 (0.376)	3.452 (0.449)	3.656 (0.513)
Gender	0.080 (0.005)	0.082 (0.007)	0.110 (0.006)	0.107 (0.008)	0.112 (0.009)	0.118 (0.010)
Age	0.033 (0.014)	0.139 (0.029)	0.044 (0.020)	0.125 (0.035)	0.157 (0.042)	0.143 (0.047)
Age Sq.	0.001 (0.000)	-0.001 (0.001)	0.001 (0.000)	-0.001 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Comuna Wages	0.007 (0.001)	0.005 (0.002)	0.013 (0.002)	0.012 (0.003)	0.016 (0.004)	0.013 (0.004)
Productivity Factor	0.570 (0.003)	0.571 (0.004)	0.526 (0.003)	0.536 (0.005)	0.540 (0.006)	0.530 (0.006)
Ability Factor	0.135 (0.004)	0.123 (0.005)	0.146 (0.004)	0.151 (0.006)	0.128 (0.006)	0.141 (0.007)
Precision	84.410 (2.746)	74.202 (3.361)	17.786 (0.184)	19.583 (0.445)	18.196 (0.478)	20.582 (0.578)
Sample Size	37,089					

Note: We show the estimated parameters of the earnings process (equation 6). We obtain these estimates by simulating 500 values of parameters using our estimated posterior. The ‘Sample Size’ row denotes the number of individuals included in our data and used to estimate the model. 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.

F Decomposing the Treatment Effects

Our model provides an opportunity to examine the extent to which selection-into-training accounts for differences in reduced form estimates and structural parameters. For instance, we can recover the importance of selection bias using the following decomposition:

$$\underbrace{E[Y_i(h_t)|D_i(h_t = 1)] - [Y_i(h_t)|D_i(h_t = 0)]}_{\text{Observed Difference}} = \underbrace{E[Y_i(h_t, 1) - Y_i(h_t, 0)|D_i(h_t = 1)]}_{\text{Treatment on the Treated}} + \underbrace{E[Y_i(h_t, 0)|D_i = 1] - E[Y_i(h_t, 0)|D_i = 0]}_{\text{Selection Bias}} \quad (\text{F.1})$$

And selection bias can be further decomposed into the importance of selection on observed variables and selection on latent factors (θ) (Heckman et al., 2018). We take advantage of our estimated model parameters to present empirical evidence regarding the importance of selection bias driven by the latent factors across treatment effect parameters in different training nodes:

Table F.1: Treatment Effect Decomposition

	Decomposition by Training Node		
	$t = 1$ (1)	$t = 2 (D1 = 0)$ (2)	$t = 2 (D1 = 1)$ (3)
Observed Differences	0.091 (0.001)***	0.127 (0.002)***	0.160 (0.003)***
Treatment on Treated	0.026 (0.000)***	0.040 (0.001)***	0.007 (0.001)***
Selection Bias (X)	0.014 (0.001)***	0.018 (0.001)***	0.044 (0.001)***
Selection Bias (θ)	0.051 (0.001)***	0.068 (0.002)***	0.109 (0.002)***

The results indicate that selection bias on unobservable characteristics plays a critical role in explaining the differences between reduced form estimates and the treatment effects recovered using estimated model parameters.

G Additional Results

Table G.1: Static Returns to Job Training (in %): Observables in Wage Equations

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	2.56 [2.53,2.60]	3.73 [3.65,3.81]	0.69 [0.53,0.85]
<i>TT</i> (percentage points)	2.57 [2.49,2.65]	3.82 [3.59,4.05]	0.77 [0.65,0.89]
$Pr(ATE < 0) \times 100$	43.63 [43.52,43.74]	45.47 [45.35,45.60]	49.15 [48.90,49.39]
$Pr(TT < 0) \times 100$	43.64 [43.39,43.88]	45.37 [45.08,45.67]	49.22 [48.84,49.60]

Notes: Table G.1 presents the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters. This version of the model includes observed characteristics in the wage equations. We show 95% confidence intervals in brackets.

Table G.2: Dynamic Returns to First-Period Job Training (in %): Observables in Wage Equations

	DATE	DTT	DTUT
Direct effect (short-term)	1.61 (0.01) [38%]	1.61 (0.02) [38%]	1.61 (0.01) [38%]
Direct effect (medium-term)	2.64 (0.02) [62%]	2.65 (0.05) [63%]	2.63 (0.02) [62%]
Continuation value	0.00 (0.02) [0%]	-0.07 (0.04) [-2%]	0.02 (0.02) [0%]
Total	4.25 (0.02)	4.19 (0.05)	4.26 (0.03)

Notes: Table G.2 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. This version of the model includes observed characteristics in the wage equations. We present DATE, DTT, and DTUT as percentage of mean baseline of the present value of earnings ($E[\tilde{Y}_i(0)]$). We present standard errors in parenthesis and the percentage contribution of each term in brackets.

Table G.3: Static Returns to Job Training (in %): Share of Negative Returns, No Variation from ε

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
ATE (percentage points)	2.54 [2.50,2.58]	3.99 [3.91,4.08]	0.78 [0.62,0.94]
TT (percentage points)	2.58 [2.50,2.66]	4.07 [3.84,4.30]	0.81 [0.69,0.93]
$Pr(ATE < 0) \times 100$	15.22 [15.14,15.31]	4.78 [4.73,4.84]	37.93 [37.69,38.17]
$Pr(TT < 0) \times 100$	15.30 [15.13,15.48]	4.64 [4.52,4.77]	39.79 [39.42,40.16]

Notes: Table G.3 presents the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters. We present the share of workers who experience negative returns to training across training histories and over time, which does not include incorporate idiosyncratic shocks to potential wages for its calculation. We show 95% confidence intervals in brackets.

H Returns to Job Training: Intensive-Margin and Long-Term Effects

Table H.1: Reduced-Form Estimates: Intensive-Margin Returns to Training

	(1)	(2)	(3)	(4)
Hours Trained (1-20)	0.113*** (0.006)	0.107*** (0.005)	0.051*** (0.004)	0.005* (0.003)
Hours Trained (> 20)	0.329*** (0.006)	0.249*** (0.005)	0.099*** (0.004)	0.012*** (0.003)
OLS	X			
OLS + controls		X		
OLS + initial wage			X	
Individual FE				X
Observations ($N \times T$)	74178	74178	74178	74178

Notes: Table H.1 estimates the reduced-form returns to short-duration (1-20 hours) and ‘long’-duration (> 20 hour) courses. 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. Column (1) presents OLS regressions without control variables. Column (2) includes PSU test scores, high school GPA, a gender dummy, age, and age squared. Column (3) includes the same control variables along with the first monthly salary observed for each worker. Column (4) computes the first differences estimator. p -values are in parenthesis, where * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

Table H.2: Static Returns to Job Training: Courses > 20 Hours

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	3.85 [3.80,3.89]	5.35 [5.27,5.44]	3.54 [3.28,3.80]
<i>TT</i> (percentage points)	3.54 [3.40,3.68]	5.54 [5.21,5.87]	2.82 [2.65,2.98]
$Pr(ATE < 0) \times 100$	41.45 [41.33,41.58]	43.39 [43.25,43.52]	45.71 [45.31,46.10]
$Pr(TT < 0) \times 100$	42.05 [41.66,42.44]	43.43 [42.99,43.86]	47.34 [46.63,48.05]

Notes: Table H.2 presents the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters along with the share of workers who experience negative returns to training across training histories and over time following from an alternative version of the model in which we exclude workers who took training courses lasting between 1-20 hours. As such, the estimated returns recover the effects of training participation in courses lasting longer than 20 hours.

Table H.3: Dynamic Returns to First-Period Training: Courses > 20 Hours

	DATE	DTT	DTUT
Direct effect (short-term)	2.37 (0.01) [37%]	2.21 (0.04) [36%]	2.39 (0.01) [37%]
Direct effect (medium-term)	3.74 (0.02) [58%]	3.74 (0.07) [61%]	3.74 (0.03) [58%]
Continuation value	0.29 (0.01) [5%]	0.20 (0.05) [3%]	0.30 (0.01) [5%]
Total	6.41 (0.03)	6.15 (0.08)	6.44 (0.03)

Notes: Table H.3 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 following from an alternative version of the model in which we exclude workers who took training courses lasting between 1-20 hours. As such, the estimated returns recover the effects of training participation in courses lasting longer than 20 hours.

Table H.4: Static Returns to Job Training: Long-Term' Earnings

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	2.43 [2.39,2.48]	4.57 [4.47,4.66]	1.22 [1.05,1.39]
<i>TT</i> (percentage points)	2.46 [2.37,2.55]	4.73 [4.48,4.98]	1.40 [1.27,1.52]
$Pr(ATE < 0) \times 100$	44.70 [44.59,44.82]	44.88 [44.75,45.01]	48.66 [48.41,48.91]
$Pr(TT < 0) \times 100$	44.72 [44.47,44.96]	44.71 [44.41,45.00]	48.64 [48.27,49.01]

Notes: Table H.4 presents the estimated Average Treatment Effects (ATE) and Treatment on the Treated (TT) parameters along with the share of workers who experience negative returns to training across training histories and over time following from an alternative version of the model in which we examine the returns to training on 'long-term' earnings following the second training stint, as defined above.

Table H.5: Dynamic Returns to First-Period Training: 'Long-Term' Earnings

	DATE	DTT	DTUT
Direct effect (short-term)	0.97 (0.01) [17%]	0.98 (0.02) [17%]	0.96 (0.01) [16%]
Direct effect (medium-term)	4.62 (0.03) [79%]	4.57 (0.07) [79%]	4.64 (0.04) [79%]
Continuation value	0.28 (0.02) [5%]	0.25 (0.05) [4%]	0.28 (0.03) [5%]
Total	5.87 (0.03)	5.79 (0.07)	5.89 (0.04)

Table H.5 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 following from an alternative version of the model in which we examine the returns to training on 'long-term' earnings following the second training stint, as defined above.

I Returns to Job Training: Mixture of Three Normals

Table I.1: Static Returns to Training: Mixture of 3 Normals

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	2.79 [2.75,2.83]	4.07 [3.99,4.16]	0.72 [0.56,0.88]
<i>TT</i> (percentage points)	2.81 [2.73,2.89]	4.07 [3.84,4.30]	0.86 [0.74,0.98]


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 I.2: Dynamic Returns to Training: Mixture of 3 Normals

	DATE	DTT	DTUT
Direct effect (short-term)	1.71 (0.01) [37%]	1.73 (0.02) [38%]	1.71 (0.01) [37%]
Direct effect (medium-term)	2.91 (0.02) [63%]	2.90 (0.05) [63%]	2.92 (0.02) [63%]
Continuation value	0.00 (0.02) [0%]	-0.05 (0.04) [-1%]	0.02 (0.02) [0%]
Total	4.63 (0.02)	4.58 (0.05)	4.65 (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.

Figure I.1: Goodness of Fit: Training Decisions



Note: In Figure I.1, 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 I.2: Distribution of Latent Factors



Note: In Figure 3, we show the estimated density of the latent factors. 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 three normal distributions. The first panel presents the density of the latent ability factor (θ_1), whereas the second panel shows the corresponding density for the latent productivity factor (θ_2).

J Returns to Job Training: Stayer Sample

Table J.1: Static Returns to Training: Stayer Sample

Treatment effect	$t = 1$	$t = 2 (D_1 = 0)$	$t = 2 (D_1 = 1)$
<i>ATE</i> (percentage points)	3.72 [3.68,3.76]	3.16 [3.08,3.25]	1.04 [0.87,1.20]
<i>TT</i> (percentage points)	3.74 [3.66,3.82]	3.15 [2.90,3.39]	1.03 [0.91,1.15]

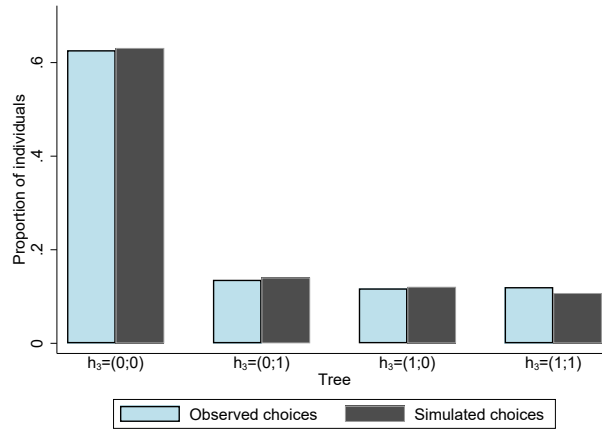
Notes: We report estimates of the Average Treatment Effects (ATE), Treatment on the Treated (TT), and likelihood of negative treatment effects across three training nodes. The first column show treatment effects after individuals make their first choice ($t = 1$). The second and third columns show estimates 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 J.2: Dynamic Returns to Training: Stayer Sample

	DATE	DTT	DTUT
Direct effect (short-term)	2.14 (0.01) [41%]	2.14 (0.02) [41%]	2.14 (0.01) [41%]
Direct effect (medium-term)	3.00 (0.02) [58%]	3.01 (0.04) [58%]	2.99 (0.02) [57%]
Continuation value	0.07 (0.02) [1%]	0.02 (0.04) [0%]	0.09 (0.02) [2%]
Total	5.21 (0.02)	5.17 (0.05)	5.22 (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.

Figure J.1: Goodness of Fit: Training Decisions



Note: In Figure J.1, 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.