

NBER WORKING PAPER SERIES

HOROWITZ-MANSKI-LEE BOUNDS WITH MULTILAYERED SAMPLE SELECTION

Kory Kroft
Ismael Mourifié
Atom Vayalinkal

Working Paper 32952
<http://www.nber.org/papers/w32952>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024, Revised February 2025

This present version of the paper is February 5, 2025. We thank David Lee and seminar audiences at Brown University, Carnegie Mellon University, National University of Singapore, New York University, Princeton University, University of Toronto, as well as the Rochester Labor Economics and Southern Economic Association conferences. We thank Stephen Claassen for providing outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Kory Kroft, Ismael Mourifié, and Atom Vayalinkal. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Horowitz-Manski-Lee Bounds With Multilayered Sample Selection
Kory Kroft, Ismael Mourifié, and Atom Vayalinkal
NBER Working Paper No. 32952
September 2024, Revised February 2025
JEL No. C01, C21, C26, H0, H50, J0, J24, J32

ABSTRACT

This paper investigates the causal effect of job training on wage rates in the presence of firm heterogeneity. When training affects the sorting of workers to firms, sample selection is no longer binary but is “multilayered”. This paper extends the canonical Heckman (1979) sample selection model – which assumes selection is binary – to a setting where it is multilayered. In this setting Lee bounds set identifies a total effect that combines a weighted-average of the causal effect of job training on wage rates across firms with a weighted-average of the contrast in wages between different firms for a fixed level of training. Thus, Lee bounds set identifies a policy-relevant estimand only when firms pay homogeneous wages and/or when job training does not affect worker sorting across firms. We derive analytic expressions for sharp bounds for the causal effect of job training on wage rates at each firm that leverage information on firm-specific wages. We illustrate our partial identification approach with two empirical applications to job training experiments. Our estimates demonstrate that even when conventional Lee bounds are strictly positive, our within-firm bounds can be tight around 0, showing that the canonical Lee bounds may capture only a pure sorting effect of job training.

Kory Kroft
Department of Economics
University of Toronto
150 St. George Street
Toronto, ON M5S 3G7
and NBER
kory.kroft@utoronto.ca

Atom Vayalinkal
Department of Economics
University of Toronto
150 St. George Street
Toronto, ON M5S 2E9
atom.vayalinkal@mail.utoronto.ca

Ismael Mourifié
Department of Economics
Washington University in St. Louis
One Brookings Drive
St. Louis, MO 63130-4899
and NBER
ismaelyacoub@gmail.com

1. INTRODUCTION

Governments allocate substantial funds to job training programs that are designed to improve worker skills. The United States (U.S.) federal government spends roughly 19 billion U.S. dollars annually on employment and training programs.¹ Federal agencies administer roughly 40 employment and training programs to assist job seekers in gaining employment. Against this backdrop, there is an ongoing debate about the appropriate level of spending on such programs. Advocates argue that they help close the “skills gap” and address worker shortages, while critics argue that they are ineffective and socially wasteful.

At the center of the debate is the long-standing question of whether job training has a causal effect on the labor market outcomes of participants. This question has garnered significant interest from both academics and policymakers and knowing the answer is essential for determining whether to continue spending on training programs. Substantial progress has been made to answer this question as a result of randomized evaluations of job training programs. These evaluations have been the subject of several comprehensive meta-analyses (see Heckman et al. 1999 and Card et al. 2010, 2018).

To date, most training program evaluations have focused on total earnings as the main outcome of analysis. While the earnings impact is surely important for answering some questions, for others it is important to narrow the focus. Earnings naturally reflect both labor supply decisions (employment and hours margin) and wage rates. To better understand whether job training increases worker skills and welfare, standard economic models show that it is important to focus on the latter.^{2,3} However, identifying the causal effect of job training on wage rates is empirically challenging due to the well-known sample selection problem (Heckman 1979). This arises since a researcher only observes wages of the employed, and the likelihood of employment can itself be impacted by job training.

In a seminal contribution, Lee (2009) showed that the causal effect of job training on wages for the always-employed can partially identify under the Imbens and Angrist (1994) (“IA”) monotonicity assumption. Lee’s key insight was to reduce the partial identification

¹See the Council of Economic Advisers 2019 report: “Government Employment and Training Programs: Assessing the Evidence on their Performance” (The Council of Economic Advisers 2019).

²Labor supply decisions could also be impacted by job training via an increase in human capital. In particular, workers with higher skills are more likely to be offered – and accept – better paying jobs.

³Hendren and Sprung-Keyser (2020) measure the willingness to pay for job training using the treatment effect on total earnings. This assumes that all increases in earnings come from returns to human capital (higher wage rate), not from higher levels of labor supply.

problem in this framework to the one considered by Horowitz and Manski (1995) – the problem of finding sharp bounds for the mean of an unobserved potential outcome that is a component of an observed mixing distribution with set-identified mixing probabilities. IA’s monotonicity condition delivers point identification of both the mixing weight (i.e., the share of always-employed) and the mean of the “untrained” potential outcome for the always-employed. Therefore, to obtain sharp bounds on the causal effect of job training on wages, one needs only to find sharp bounds on the mean of the “trained” potential outcome for the always-employed. Lee’s bounding approach has become influential in empirical research.⁴

In developing his approach, Lee focused on the potential for job training to affect labor supply along the extensive margin (work vs. no work). Although it is important to know whether training increases employment, a fundamental question for policymakers is whether job training improves labor market outcomes by increasing job quality. If this is the case, training may increase both the probability individuals are employed, as well as the probability individuals are employed at “better jobs”, both of which can increase earnings. For example, this is a key feature of President Biden’s workforce training initiative, the “American Rescue Plan’s Good Jobs Challenge”, which prioritizes job quality and is designed to ensure that workers can access good jobs.

Despite the emphasis on job quality by policymakers, the academic literature on job training programs has mostly ignored firms. In the standard competitive model considered by Lee and most of the training literature, the firm for which an individual works does not matter for wages. However, there is growing empirical evidence that demonstrates the importance of firms for wage determination.⁵ This research highlights the importance of having a “good job” which can be interpreted as working at a “good firm” that offers a higher wage for all its employees. Thus, an open question is whether job training raises earnings by moving participants to higher-paying firms.

⁴In Appendix F, we describe results from an informal survey we conducted. We counted 56 papers published in ‘top 5’ general interest economic journals – the American Economic Review, Econometrica, the Journal of Political Economy, the Quarterly Journal of Economics and the Review of Economic Studies – that cited Lee (2009) or earlier working paper versions. The purpose of our survey was to determine whether the nature of sample selection in these papers was binary or multilayered. Of these, 42 empirically implemented Lee bounds to address sample selection, and 7 of them featured multilayered selection. To apply Lee bounds in these settings, the researchers collapsed the sample selection problem to a single dimension.

⁵See, for example, Abowd et al. (1999), Card et al. (2013), Song et al. (2019), Bonhomme et al. (2019) and Bonhomme et al. (2023).

The presence of firm heterogeneity and the potential for worker sorting raises several new questions of interest. First, what estimand do “Lee bounds” partially identify when there is firm heterogeneity in wages? This question cannot be addressed with the canonical Heckman (1979) sample selection model since this assumes that sample selection is binary, i.e., job training can increase employment but has no effect on sorting to firms. The first contribution of this paper is to extend the standard sample selection model to a setting where selection is multilayered, and to show that the conventional Lee bounds set identifies (for the always-employed population) a total effect that combines a weighted average of the causal effect of job training on wage rates across firms (we label this the “within-firm effect”) with a weighted average of the contrast in wages between different firms for a fixed level of training (we label this the “sorting effect”).

A natural question is whether it is possible to separate the within-firm wage effect of job training from the sorting effect in the presence of heterogeneous firms. There are several reasons why one would want to separately identify these effects. First, some features of job training programs affect sorting (job search assistance), whereas others affect skill acquisition (classroom and vocational training). Thus, the decomposition could potentially highlight which investments – job search assistance or classroom training – are effective for raising wages and thus improve targeting. Second, the within-firm wage effect is arguably better able to shed light on the causal effect of *employer-sponsored* job training. Third, for a welfare analysis of job training programs, it is important to focus on the direct wage effects of job training, since labor supply effects have second-order effects on utility (via the envelope theorem) (see Hendren and Sprung-Keyser 2020).⁶ Thus, we argue that the within-firm wage effect is the more relevant causal effect for evaluating welfare effects.

The second contribution of this paper is to derive sharp bounds on the within-firm wage effect. Our bounding approach proceeds in two steps. In the first step, we derive sharp closed-form bounds on the *response type* probabilities.⁷ In deriving these bounds, we exploit a unique feature of our setting, which is that (unlike in the traditional instrumental variables framework), the exclusion restriction does not hold, since job training can have a direct causal effect on the outcome (wages). We show that this feature implies that the distribution of response types does not depend on the outcome (wage) distribution and allows us to derive

⁶This logic requires that the government is increasing spending on job training by a sufficiently small amount.

⁷The *response type* represents the pair of firms that an individual would choose to work at if she were externally assigned to the control group or the treatment group, respectively.

closed-form bounds on the distribution of response types.⁸ The second step provides closed-form bounds on the treatment effects, as a function of the sharp bounds on the response types derived in the first step. This step involves extending the Horowitz and Manski (1995) approach (which involves a single-equation mixture model with two components) to our setting which involves two mixture model equations with unknown weights that are interdependent across the equations. Importantly, we show that while this two-step approach provides an easy and tractable way to construct closed-form bounds, it does not entail any loss of information and provides sharp bounds. We also consider a set of additional restrictions on response types and show that they naturally lead to tighter bounds on the treatment effects of interest. Using a numerical illustration, we demonstrate that our bounds can be informative enough to potentially separate cases with a positive within-firm effect from those with no within-firm effect, even when both cases lead to strictly positive Lee bounds.

We consider two empirical applications. Our first application is based on the randomized evaluation of Job Corps following Lee (2009). We classify firms into observable firm types taking advantage of the fact that in the publicly available survey data there are direct measures of firm amenities, such as the availability of health insurance, paid vacation and retirement, or pension benefits. We show that, on average, firms that offer amenities pay higher wages than firms that do not. Moreover, the wage distribution for amenity-offering firms stochastically dominates the wage distribution for non-amenity-offering, within both the treatment and control groups. We then go on to show that being randomly assigned to Job Corps led individuals to work at firms with better job amenities compared to the control group. This combined evidence suggests that selection is multilayered and motivates our implementation of sharp bounds to evaluate Job Corps. We begin by replicating the findings of Lee (2009). Our estimates reveal that even in the cases where conventional Lee bounds are strictly positive ($[0.047, 0.048]$, 90 weeks after training), our multilayered bounds for the within-firm wage effect (which hold the sorting effect constant) include 0. This suggests that Lee bounds may capture a pure sorting effect of job training rather than a direct human capital effect. Indeed, under additional plausible assumptions on the distribution of response-types, our bounds on the within-firm wage effect become tight but continue to contain 0.

⁸While we derive closed-form bounds on the distribution of response types, we show that one can obtain them equivalently using a linear programming approach.

Our second empirical application is based on the WorkAdvance experiment recently examined in Katz et al. (2022). This is a sectoral employment program that targets high-quality jobs in specific industries with strong labor demand. We classify firms based on whether a firm operates in a WorkAdvance targeted sector. Similar to Job Corps, wages are higher at target sector firms and random assignment to job training increases sorting to these firms. We again document a case where Lee bounds are strictly positive and tight, but our within-firm wage effect includes 0.

The remainder of the paper is organized as follows. Section 2 considers the multilayered sample selection problem in both a parametric model along the lines of Heckman (1979) and a more general treatment effects framework; it defines the key causal estimands of interest. Section 3 considers the causal interpretation of Lee bounds in the presence of multilayered sample selection and presents the general decomposition. Section 4 derives the sharp bounds of a large class of parameters of interest in the multilayered sample selection model. Section 5 presents empirical applications that implement the sharp bounds for Job Corps and WorkAdvance. The last Section concludes. The Appendices contain simulations of the model assuming that there are two types of firms, and presents all the proofs for the paper.

Related Literature. Our paper builds on and contributes to the following literature. First, there is a large literature on active labor market programs that is reviewed in Heckman et al. (1999) and Card et al. (2010, 2018). Our contribution to this literature is to examine whether, and to what extent, worker sorting to firms affects the wage impacts of job training. To our knowledge, there are only a few empirical analyses that have examined the impact of training on worker sorting to firms. Andersson et al. (2022) find suggestive evidence of a positive impact of training on firm characteristics, as well as effects on industry of employment. Another related study is Katz et al. (2022), who evaluate sector-based training programs. Examining evidence from randomized evaluations of programs that combine upfront screening, occupational and soft skills training, wrap-around services, and targeted low-wage workers, Katz et al. (2022) find substantial and persistent earnings gains after training. With regard to mechanisms, they interpret the earnings gain as driven in part by the sorting of workers to higher-paying industries and occupations. They do not, however, provide a framework for isolating the sorting effect as a causal mechanism. Finally, Schochet et al. (2008) evaluate the impact of Job Corps on the sorting of workers to jobs with different amenities, such as the availability of health insurance, and retirement or pension benefits, and report positive impacts. However, they do not disentangle the effects on these

characteristics for those who would be employed in any case from a selection effect that comes from the impact of Job Corps on employment.

Second, our paper relates to the literature that has documented firm heterogeneity in wages. Firms have been shown to be important for wage inequality (Abowd et al. 1999), the cyclical nature of wages and early career progression (Card et al. 2013), the earnings losses of displaced workers (Lachowska et al. 2020; Schmieder et al. 2023), and gender (Card et al. 2016) and racial wage gaps (Gerard et al. 2021). Our contribution to this literature is to examine the role of firms in understanding the wage effect of job training.

Third, our paper relates to econometric approaches that address the sample selection problem. The Heckman (1979) sample selection model has been extended in various dimensions. First, a series of papers, including Gallant and Nychka (1987), Newey et al. (1990), and Ahn and Powell (1993), propose estimation and inference methods that relax the normality assumption imposed by Heckman (1979); see Li and Racine (2007, Chapter 10) for a review of such extensions. Second, Lee (2009) extends Heckman (1979) by relaxing the exclusion restriction of instrumental variables and derives bounds on the parameters of interest. Honoré and Hu (2020) study a semiparametric version of Lee’s model. Additionally, Semenova (2020) and Olma (2021) propose various approaches for inference on Lee’s bounds conditional on (potentially continuous) covariates. To our knowledge, this paper represents the first attempt to extend the seminal Heckman (1979) sample selection model to multilayered settings. Although the focus of our paper is primarily on firms, which we consider to be the main layer of interest, our analysis can be extended in various directions. For example, one could consider occupation as a layer and examine the returns to occupation while controlling for sorting, similar to the approach taken by Gottschalk et al. (2014).

Finally, one can view the firm as a “mediator” in the context of the literature on mediation analysis (see, for example, Robins and Greenland (1992) and Pearl (2001)). Traditionally, most of this literature ignores sample selection where the outcome is not observed at some mediator values. For example, recently Kwon and Roth (2024) developed a method to test for the presence of a mediator but abstracts from sample selection. A rare exception is Zuo et al. (2022), who consider identification of direct and indirect effects within a mediation analysis framework, when both the outcome and the mediator are missing. They focus on point identification under various assumptions, including the abstract and non-falsifiable assumption of completeness; however, many of their assumptions do not apply to

our setting.⁹ For instance, in our setting, whether the wage is observed can depend on wage rate offered by firms (the mediator), and this situation is ruled out by their assumptions. Our paper therefore complements Zuo et al. (2022) by establishing partial identification of the direct and indirect effects without imposing completeness. Our maintained assumptions are transparent and directly apply to the primitives of our model. Moreover, our approach accommodates an endogenous mediator and allows the outcome to be missing non-randomly, even conditional on covariates.

2. ANALYTICAL FRAMEWORK

2.1. Multilayered Sample Selection: A parametric model. Since Heckman’s seminal work in 1979, the sample selection model has been formalized as follows:

$$Y = \begin{cases} \alpha Z + \beta X' + U, & \text{if } D = 1, \\ \text{unobserved}, & \text{if } D = 0, \end{cases} \quad (2.1)$$

$$D = 1\{\eta Z + \theta X' + V > 0\}, \quad (2.2)$$

where D captures the binary sample selection model (i.e., D is equal to 1 if employed and 0 if not) and Y represents the outcome which is observed only when D is equal to 1 (i.e., Y is the observed wage when employed). The latent variables in the model are denoted as (U, V) , X is a vector of observed exogenous covariates, and $Z \in \{0, 1\}$ is a binary variable that respects the conditional independence assumption: $(U, V) \perp Z | X$ (i.e., job training, Z , is randomly assigned). In the Heckman sample selection model, identification of the parameters in the outcome equation requires at least one variable that is independent of the latent variables but is excluded from the outcome equation. In model (2.1, 2.2) this is equivalent to assuming α to be equal to 0, implying that Z is excluded from the outcome equation. In this context, Z becomes a valid instrumental variable to consistently estimate β , satisfying both the independence and exclusion restriction conditions.

However, as highlighted in Lee (2009) and recognized more generally, the exclusion restriction can be violated in certain cases implying that $\alpha \neq 0$. Moreover, in such cases, α is potentially a parameter of primary interest. For instance, in the job training example, participating in training could boost an individual’s human capital and directly affect their wage rate. Consequently, α can be interpreted as the causal effect of job training on the

⁹For an in-depth review of completeness, see D’Haultfoeuille (2011), and also Canay et al. (2013).

wage rate, which is the key parameter studied by Lee (2009). The primary methodological contribution of Lee is to provide a method that allows researchers to partially identify α in the presence of sample selection (i.e., when training can affect labor supply through η).

A key assumption in the parametric model above is that the sample selection problem is binary: individuals are either employed or unemployed. Yet if job training affects not only whether an individual works but also which firm they work at, the sample selection problem becomes multilayered. We now generalize the seminal Heckman sample selection model (2.1, 2.2) to allow for a richer model of labor supply where individuals choose layers, i.e., firms. We refer to this extended model as the “*parametric multilayered selection model*”:

$$Y = \begin{cases} \alpha_K Z + \beta_K X' + U_K, & \text{if } D = K, \\ \vdots & \vdots \\ \alpha_1 Z + \beta_1 X' + U_1, & \text{if } D = 1, \\ \text{unobserved}, & \text{if } D = 0, \end{cases} \quad (2.3)$$

$$D = \arg \max_{d \in \{0, 1, \dots, K\}} \{\eta_d Z + \theta_d X' + V_d\} \quad (2.4)$$

where $\eta_0 Z + \theta_0 X' + V_0 = 0$. In this model, each layer (D) represents a distinct firm, with corresponding parameters α_d , β_d , and latent variable U_d . Expected utility for a given firm d is given by $\eta_d Z + \theta_d X' + V_d$. The utility of the outside option (i.e., unemployment) is $\eta_0 Z + \theta_0 X' + V_0 = 0$. The worker selects the firm with the highest expected utility.

In the parametric multilayered sample selection model, α_d is the causal effect of job training on the wage rate at firm d . We refer to this causal effect as the “within-firm effect” for layer d . The vector (η_1, \dots, η_K) includes parameters that reflect the causal effect of job training on firm choice. In the next section, we demonstrate that Lee bounds combine both the within-firm effects $(\alpha_1, \dots, \alpha_K)$ with the sorting effects summarized by (η_1, \dots, η_K) . As discussed in the Introduction, it is interesting to separately identify these causal effects. For example, it is straightforward to show that the change in utility associated with a small increase in spending on job training is captured by the mechanical increase in wages. The sorting effect has only a second-order effect on worker utility due to the envelope theorem.

2.2. Multilayered Sample Selection: Generalized version using the potential outcome model. Let (Ω, \mathcal{F}, P) be a probability space, where we interpret Ω as the population of interest, and $\omega \in \Omega$ as a generic individual in the population. Let $Y_{z,d}(\omega)$ be the potential outcome (i.e., potential wage) if agent ω is externally assigned to the treatment group

$z \in \{0, 1\}$ (i.e., job training) and to a specific layer $d \in \{0, 1, \dots, K\}$, where $d = 0$ denotes the layer for which the outcome is not observed (i.e., $Y_{z,0}(\omega)$ is not observed).¹⁰ D_z denotes the potential layer that the individual selects if externally assigned to the treatment group $z \in \{0, 1\}$. We denote the realized outcome by $Y \in \mathcal{Y} \subseteq \mathbb{R}$ and the realized layer by $D \in \{0, 1, \dots, K\}$. Finally, let Z be the assigned treatment group and let X be a vector of covariates. We assume that (D, Z, X) is observed for every individual, but the realized outcome Y is only observed if $D \neq 0$ (i.e., if the individual is employed). This implies the following model:¹¹

$$Y = \sum_{d=1}^K [Y_{1,d}Z + Y_{0,d}(1 - Z)] 1\{D = d\}, \quad (2.5)$$

$$D = D_1Z + D_0(1 - Z), \quad (2.6)$$

along with the following conditional independence assumption.

Assumption 1 (Conditional Random Assignment). *Individuals are randomly assigned to a treatment group. $\{(Y_{z,d}, D_z) : d \in \{0, 1, \dots, K\}, z \in \{0, 1\}\} \perp Z | X$.*

The outcome equation (2.5) collapses to the outcome model of Lee (2009) when there is no heterogeneity between the potential layered outcomes, that is, $Y_{z,d} = Y_z$ for $d \in \{1, \dots, K\}$. In this case, we have:

$$\begin{aligned} Y &= \sum_{d=1}^K [Y_{1,d}Z + Y_{0,d}(1 - Z)] 1\{D = d\} = [Y_1Z + Y_0(1 - Z)] \sum_{d=1}^K 1\{D = d\} \\ &= [Y_1Z + Y_0(1 - Z)] 1\{D \neq 0\}. \end{aligned}$$

What causal interpretation should be given to Lee bounds in the presence of multilayer sample selection, where $Y_{z,d} \neq Y_{z,d'}$, for $d, d' \in \{1, \dots, K\}$? Depending on the researcher's interest, various causal estimands of interests could be defined. Before defining our parameters of interest, we show that there is a link between the causal effects in our multilayered framework and those typically considered in the mediation analysis literature. We use this link to characterize our key estimands below.

¹⁰In our empirical application, we will assume that the layer corresponds to a firm's type, where the type is constructed based on a firm's observable characteristics.

¹¹Strictly speaking, equation (2.5) implies that the outcome $Y_{z,d}(\omega) = 0$ when $D = 0$, but it should really be interpreted as $Y_{z,d}(\omega)$ is unobserved.

2.3. Direct and indirect effects in presence of sample selection. In our model, particularly in equation (2.5), there is a notable connection to the literature on mediation analysis, as discussed by Pearl (2001) and others. The graphical representation of the outcome equation in our model takes the form:¹²




FIGURE 1. DAG of causal relationships between variables in our model.

In the context of mediation analysis, where Z represents the randomized treatment, D is conceptualized as the “mediator”, and Y denotes the outcome, our model allows the treatment (job training) to influence the outcome through two channels: a direct channel and an indirect channel that passes through the mediator. W is a vector of latent unobserved variables often called “confounding variables” which simultaneously affect D and Y , making D an endogenous variable. In the parametric model, $W \equiv (U_1, \dots, U_K, V_1, \dots, V_K)$. In our framework, the mediator corresponds to the firm where the individual would be employed if they were externally assigned to job training.

In the mediation analysis literature, two categories of causal estimands have garnered attention: “direct effects” and “indirect effects”. Focusing on the former, two types of “direct effects” have been conceptualized. First, the *control direct effect* (CDE) is defined as:

$$\text{CDE}(d) \equiv \mathbb{E}[Y_{1,d} - Y_{0,d}]. \quad (2.7)$$

This captures the causal effect of job training on earnings within a specific firm d when the firm is held fixed. It is equivalent to the within-firm wage effect for layer d . In Lee’s (2009) terminology, the CDE corresponds to the causal impact of job training on the wage rate illustrated by the curved arrow in Figure 1. This parameter is the primary focus of Lee (2009).

¹²For the sake of clarity, this graph simplifies the discussion by omitting sample selection.

The CDE can vary significantly across firms (d), reflecting the potentially heterogeneous impact of job training on wages for different firms. CDEs are useful when the policymaker is primarily interested in the impact of job training on wages at a specific firm. More generally, policymakers may also be interested in understanding the overall impact of training on wages at firms that workers naturally choose when they receive training. The second type of direct effect – the “natural direct effect” (NDE) – is introduced to capture this notion:

$$\text{NDE} \equiv \mathbb{E}[Y_{1,D_1} - Y_{0,D_1}] = \sum_{d=0}^K \mathbb{E}[Y_{1,d} - Y_{0,d} | D_1 = d] \times \mathbb{P}[D_1 = d] \quad (2.8)$$

where $Y_{z,D_{z'}} \equiv \sum_{d=0}^K Y_{z,d} 1\{D_{z'} = d\}$ for $z, z' \in \{0, 1\}$, and $d \in \{0, \dots, K\}$. The expression $Y_{1,D_1}(\omega) - Y_{0,D_1}(\omega)$ represents the causal impact of job training on wages for the specific firm that the worker ω would have selected if she had been externally assigned to receive job training. The NDE is essentially the average of these individual effects.

Turning to indirect effects, the “natural indirect effect” (NIE) is defined as:

$$\begin{aligned} \text{NIE} &\equiv \mathbb{E}[Y_{0,D_1} - Y_{0,D_0}] \\ &= \sum_{d=0}^K \sum_{d'=0: d' \neq d}^K \mathbb{E}[Y_{0,d} - Y_{0,d'} | D_0 = d', D_1 = d] \times \mathbb{P}[D_0 = d', D_1 = d]. \end{aligned} \quad (2.9)$$

The term $Y_{0,d} - Y_{0,d'}$ represents the wage contrast between firms d and d' in the absence of job training. However, rather than specifying the pair of firms (d, d') , we can examine this wage difference at the “natural representative” firms D_1 and D_0 , resulting in $Y_{0,D_1} - Y_{0,D_0}$. The indirect effect aims to capture the causal impact of job training on the outcome purely due to the change in firms, and can therefore be seen as the influence of job training transitioning through a change in the firm. As highlighted by Pearl (2009), the empirical relevance of the indirect effect estimand is controversial and questionable. Implementing an intervention that would suppress the direct effect of Z on Y while allowing the indirect channel through D is not realistic. However, it remains a key parameter in the mediation analysis literature.

In the context of sample selection, the outcome is observable only when $D \neq 0$. We can categorize the population into four major groups: $\Omega = \{\omega : D_0(\omega) = 0, D_1(\omega) = 0\} \cup \{\omega : D_0(\omega) > 0, D_1(\omega) = 0\} \cup \{\omega : D_0(\omega) = 0, D_1(\omega) > 0\} \cup \{\omega : D_0(\omega) > 0, D_1(\omega) > 0\}$. For the first group we never observe outcomes, regardless of training status. For the second group only outcomes under job training are never observed, whereas for the third group only outcomes when not assigned to job training are never observed. If we are unwilling

to assume that the outcome is missing at random (or selection on observables only) or impose parametric assumptions, the observed data cannot provide information on the causal effect of job training for individuals belonging to those groups. We refrain from imposing such stringent restrictions and focus solely on the causal effects for the final group, the subpopulation $\{\omega : D_0(\omega) > 0, D_1(\omega) > 0\}$.

It is useful to further divide our population of interest $\{\omega : D_0(\omega) > 0, D_1(\omega) > 0\}$ into finer groups, which we label *response types*, i.e. $\{\omega : D_0(\omega) > 0, D_1(\omega) > 0\} = \cup_{\{d, d' \in \{1, \dots, K\}\}} \{\omega \in \Omega : D_1(\omega) = d, D_0(\omega) = d'\}$.¹³ Response types are defined by the pair of firms that individual ω would choose to work for if externally assigned to the control group or the treatment group. Formally, the response type is defined as the random variable $T = (D_0, D_1)$ and \mathcal{T} represents its support.

We now introduce two pivotal parameters, the Local Controlled Direct Effect (LCDE) and the Local Controlled Indirect Effect (LCIE):

$$\text{LCDE}(d|t) = \mathbb{E}[Y_{1,d} - Y_{0,d}|T = t], d \in \{1, \dots, K\}, \text{ and } t \in \mathcal{T} \quad (2.10)$$

and

$$\text{LCIE}(z, d, d'|t) = \mathbb{E}[Y_{z,d} - Y_{z,d'}|T = t], d \in \{1, \dots, K\}, \text{ and } t \in \mathcal{T} \quad (2.11)$$

The individual CDEs may vary between individuals, i.e. $Y_{0,d}(\omega) - Y_{0,d'}(\omega) \neq Y_{0,d}(\omega') - Y_{0,d'}(\omega')$ for $\omega \neq \omega'$. By considering the LCDE, we allow for heterogeneity of the CDE across response types. Note that in certain instances, a specific LCDE may be more policy-relevant than the CDE itself. Both the LCDE and CDE exhibit their own policy relevance, akin to the extensive debate in the instrumental variable (IV) literature regarding the empirical relevance between the average treatment effect (ATE) and local ATE (LATE). This analogy extends to the LCIE.

We now show that the “sample selection” versions of the CDE, NDE, and NIE can be perceived as a weighted average of $\text{LCDE}(d|t)$ or $\text{LCIE}(z, d, d'|t)$. By “sample selection” version, we mean that the effects are defined conditionally on being always employed, i.e. $\{\omega : D_0(\omega) > 0, D_1(\omega) > 0\}$.

¹³See Heckman and Pinto (2018) for a more detailed discussion on the advantages of such a partition.

$$\begin{aligned} \mathbb{E}[Y_{1,d} - Y_{0,d} | D_0 > 0, D_1 > 0] &= \sum_{l=1}^K \sum_{l'=1}^K \text{LCDE}(d|l, l') \times \mathbb{P}[T = (l, l') | D_0 > 0, D_1 > 0], \\ \mathbb{E}[Y_{1,D_1} - Y_{0,D_1} | D_0 > 0, D_1 > 0] &= \sum_{d=1}^K \sum_{d'=1}^K \text{LCDE}(d|d', d) \times \mathbb{P}[T = (d', d) | D_0 > 0, D_1 > 0], \\ \mathbb{E}[Y_{0,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] &= \sum_{d=1}^K \sum_{d'=1: d \neq d'}^K \text{LCIE}(0, d, d' | d', d) \times \mathbb{P}[T = (d', d) | D_0 > 0, D_1 > 0]. \end{aligned}$$

As illustrated, $\text{LCDE}(d|t)$ or $\text{LCIE}(z, d, d'|t)$ represent more primitive parameters compared to CDE, NDE, and NIE. This paper will focus in particular on identifying the $\text{LCDE}(d|t)$. It should be noted that, in the absence of individual heterogeneity, whenever $Y_{z,d}(\omega) = Y_{z,d}(\omega')$ for $\omega \neq \omega'$, we have $\text{LCDE}(d|t) = \text{LCDE}(d) = \alpha_d$, as in the parametric version of the model. Moreover, since the outcome is never observed when $D = 0$, hereafter we use the following notation $Y_{z,D_{z'}} \equiv \sum_{d=1}^K Y_{z,d} 1\{D_{z'} = d\}$ for $z, z' \in \{0, 1\}$.

Remark 1. *In Lee (2009), training assignment Z is a randomly assigned treatment that fails to satisfy the exclusion restriction required to address sample selection using standard methods. In our more general setting, while Z remains a treatment of interest, it also plays the role of being an instrument for D . In some settings, the impact of D on the outcome may be of independent interest, and our results also apply to such settings. Moreover, since our results apply even when there is no sample selection (i.e. Y is always observed, $P(D = 0) = 0$), our approach generalizes the IV model to settings where the instrument does not satisfy the exclusion restriction.*

3. THE CAUSAL INTERPRETATION OF LEE'S BOUNDS IN THE PRESENCE OF MULTILAYERED SAMPLE SELECTION

First, notice that the generalized multilayered selection model, i.e., equations (2.5, 2.6) implies the following:

$$Y = [Y_{1,D_1}Z + Y_{0,D_0}(1 - Z)], \quad (3.1)$$

$$1\{D > 0\} = 1\{D_1 > 0\}Z + 1\{D_0 > 0\}(1 - Z). \quad (3.2)$$

Additionally, Lee imposes the following monotonicity assumption:

Assumption 2 (Conditional Lee’s Monotonicity Assumption). *We impose the following restriction: $\mathbb{P}[1\{D_1 > 0\} \geq 1\{D_0 > 0\} | X] = 1$ a.s.*

This assumption means that being assigned to the treatment group can never lower employment, and this applies uniformly for all agents in the population. In our general framework with multilayered sample selection, this assumption requires that all agents are more likely to join an employment layer when assigned to treatment. Since Lee’s monotonicity assumption is only required to hold conditional on X , it can be modified to allow the direction of monotonicity to vary across different values of X . Such a modification does not present a challenge for our identification analysis, which holds X fixed throughout, but inference methods need to be adapted to accommodate such an assumption, especially when X is continuous. For further details on such adaptations, see Słoczyński (2020) and Semenova (2020), which provide inference methods that are valid under such assumptions.

All remaining analysis, results, and assumptions should be understood as implicitly conditioning on $X = x$ for some value x of the vector of observed covariates, X , which will generally be suppressed in notation.

Lemma 1 (Lee Bounds). *Under Assumptions 1 and 2 Lee bounds set identifies the following estimand $\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0]$:*

$$\underline{\theta}^\ell \leq \mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] \leq \bar{\theta}^\ell \quad (3.3)$$

where

(i) *For continuous outcome:*

$$\underline{\theta}^\ell \equiv \mathbb{E}[Y | D > 0, Z = 1, Y \leq F_{Y|D>0,Z=1}^{-1}(p)] - \mathbb{E}[Y | D > 0, Z = 0], \quad (3.4)$$

$$\bar{\theta}^\ell \equiv \mathbb{E}[Y | D > 0, Z = 1, Y \geq F_{Y|D>0,Z=1}^{-1}(1-p)] - \mathbb{E}[Y | D > 0, Z = 0], \quad (3.5)$$

(ii) *For binary outcome:*

$$\underline{\theta}^\ell \equiv \max \left\{ 0, 1 - \frac{1}{p} P[Y = 0 | D > 0, Z = 1] \right\} - \mathbb{E}[Y | D > 0, Z = 0], \quad (3.6)$$

$$\bar{\theta}^\ell \equiv \min \left\{ 1, \frac{1}{p} P[Y = 1 | D > 0, Z = 1] \right\} - \mathbb{E}[Y | D > 0, Z = 0], \quad (3.7)$$

with $F_W^{-1}(u) \equiv \inf\{w \in \mathbb{R} : \mathbb{P}(W \leq w) \geq u\}$ for $u \in [0, 1]$ and $p \equiv \frac{\mathbb{P}(D>0|Z=0)}{\mathbb{P}(D>0|Z=1)}$.

Lemma 1 shows that in the presence of heterogeneous firms, Lee’s identification approach bounds the estimand $\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0]$. What is the causal interpretation of this estimand? The following lemma sheds light on this.

Lemma 2 (Decomposition). *Assuming the generalized multilayered sample selection model, we have the following decomposition:*

(i) *General decomposition:*

$$\begin{aligned} & \mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] \\ &= \underbrace{\sum_{d=1}^K \sum_{d'=1}^K LCDE(d|d', d) \times \mathbb{P}[T = (d', d) | D_0 > 0, D_1 > 0]}_{\mathbb{E}[Y_{1,D_1} - Y_{0,D_1} | D_0 > 0, D_1 > 0]} \\ &+ \underbrace{\sum_{d=1}^K \sum_{d'=1:d \neq d'}^K LCIE(0, d, d'|d', d) \times \mathbb{P}[T = (d', d) | D_0 > 0, D_1 > 0]}_{\mathbb{E}[Y_{0,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0]} \quad (3.8) \end{aligned}$$

(ii) *No mediation effect (No firm-specific wage rate, i.e., $Y_{z,d} = Y_{z,\bullet}$) or no sorting across firms, i.e., $\mathbb{P}[T = (d', d) | D_0 > 0, D_1 > 0] = 0$ for $d \neq d'$.*

$$\begin{aligned} & \mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] \\ &= \sum_{d=1}^K \sum_{d'=1}^K \mathbb{E}[Y_{1,\bullet} - Y_{0,\bullet} | D_1 = d, D_0 = d'] \mathbb{P}(D_1 = d, D_0 = d' | D_0 > 0, D_1 > 0) \\ &= \mathbb{E}[Y_{1,\bullet} - Y_{0,\bullet} | D_0 > 0, D_1 > 0] \quad (3.9) \end{aligned}$$

(iii) *No direct effect (i.e., $Y_{z,d} = Y_{\bullet,d}$).*

$$\begin{aligned} & \mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] \\ &= \sum_{d=1}^K \sum_{d'=1:d \neq d'}^K \mathbb{E}[Y_{\bullet,d} - Y_{\bullet,d'} | D_0 = d', D_1 = d] \times \mathbb{P}[D_0 = d', D_1 = d | D_0 > 0, D_1 > 0] \\ &= \mathbb{E}[Y_{\bullet,D_1} - Y_{\bullet,D_0} | D_0 > 0, D_1 > 0] \quad (3.10) \end{aligned}$$

Lemma 2 (i) shows that in the presence of firm heterogeneity, Lee’s partial identification approach establishes bounds for a total effect. This total effect combines the sample selection

version of the NDE and the NIE (i.e., conditional on $D_0 > 0$ and $D_1 > 0$), with each possessing distinct interpretations. Importantly, as discussed above, the NDE and NIE do not hold the mediator (firm) D fixed. The NDE is an average of causal effects of job training at each firm weighted by fraction of workers choosing that firm under job training. The NIE is an average of causal effects of firm on wages (in the no-job training counterfactual scenario) weighted by the share of response types choosing those firms. Thus, without additional assumptions, this approach does not allow one to separately identify the CDEs (the within-firm wage effects) from the labor supply effects or sorting effects that transit through D . Lemma 2 (ii) shows that when there are no mediation effects (or no heterogeneity in wages between firms), i.e., $Y_{z,d} = Y_{z\bullet}$ as assumed in Lee (2009) and illustrated in Figure 2, the




FIGURE 2. DAG when there are no mediation effects.

NIE vanishes while the NDE reduces to the CDE which is the target parameter in Lee's framework.¹⁴ Finally, Lemma 2 (iii) reveals that in the absence of a direct effect of job training on wages (i.e., $Y_{z,d} = Y_{\bullet d}$) as depicted in Figure 3, Lee bounds capture the effect of job training on wages coming exclusively from the sorting of individuals into different firms.




FIGURE 3. DAG when there is no direct effect.

¹⁴Notice that Figures 2 and 3 are drawn for the subpopulation of always observed, i.e. $\{\omega : D_0(\omega) > 0, D_1(\omega) > 0\}$.

This shows that, in general, interpreting Lee bounds as informative about the direct wage effect of job training is problematic, unless there is clear empirical evidence of the absence of mediation effects. Unfortunately, Lee’s approach does not provide a means to assess this. Given these challenges, the next section introduces an alternative partial identification approach that is designed to overcome these limitations and aims to partially identify the true causal impact of job training on wage rates.

4. SHARP BOUNDS IN THE MULTILAYERED SAMPLE SELECTION MODEL

In this section, we develop a partial identification strategy to recover the parameters $\text{LCDE}(d|t)$ and $\text{LCIE}(z, d, d'|t)$ that will allow us to isolate the within-firm effect of job training from the sorting effect. Under Assumption 1, the response type T is independent of Z . Assumption 2 restricts the response type support. For example, under Assumption 2, $\mathbb{P}[T = (d, 0)] = 0$ for $d \in \{1, \dots, K\}$. We denote by $f_{Y_{z,d}|D,Z}(y|d', z')$ the conditional density of $Y_{z,d}$ given $\{D = d', Z = z'\}$ and assume that it is absolutely continuous with respect to a dominating measure μ on $Y_{z,d}$. We note that $f_{Y_{z,d}|D,Z}(y, d|z) \equiv f_{Y_{z,d}|D,Z}(y|d, z)\mathbb{P}(D = d|Z = z)$. For $d, d' \in \{1, \dots, K\}$ and $z \in \{0, 1\}$, and any $y \in \mathcal{Y}$ we have the following:

$$f_{Y|D=d, Z=z}(y) = f_{Y_{z,d}|D_z}(y|d) = \sum_{d'=1}^K \frac{\mathbb{P}(D_z = d, D_{1-z} = d')}{\mathbb{P}(D = d|Z = z)} \times f_{Y_{z,d}|D_z, D_{1-z}}(y|d, d') \quad (4.1)$$

where the first equality holds under Assumption 1.

More precisely, under Assumption 1, the following system of equations characterizes the empirical content of the multilayered sample selection model.

$$f_{Y, D=d|Z=1}(y) = \sum_{d'=0}^K \mathbb{P}[T = (d', d)] \times f_{Y_{1,d}|T}(y|d', d) \quad (4.2)$$

$$f_{Y, D=d|Z=0}(y) = \sum_{d'=0}^K \mathbb{P}[T = (d, d')] \times f_{Y_{0,d}|T}(y|d, d') \quad (4.3)$$

and this holds for any $d, d' \in \{1, \dots, K\}$ and $y \in \mathcal{Y}$. The left-hand side of equations (4.2) and (4.3) are observed while the individual types, i.e., $\mathbb{P}[T = (d, d')]$ and the conditional potential outcome distributions, i.e., $f_{Y_{z,d}|T}(y|d', d)$ on the right-hand side of the equations are unknown. For a given d , the number of unknown quantities $2K + 1 + 2(K + 1)|\mathcal{Y}|$ is larger than the number of equations $2|\mathcal{Y}|$. We therefore have an under-determined system

of linear equations with unknown coefficients. As such, it is only possible to set identify these parameters. The identified set of unknown parameters could naturally shrink if the researcher is willing to impose additional assumptions, such as Assumption 2. For example, under Assumption 2, $\mathbb{P}[T = (d, 0)] = 0$ for $d \in \{1, \dots, K\}$ which implies that $\mathbb{P}[T = (d, 0)]f_{Y_{z,d}|T}(y|d, 0) = 0$ for $d \in \{1, \dots, K\}$. Consequently, for a fixed d , this leads to a reduction of $|\mathcal{Y}| + 1$ in the total number of unknown parameters while keeping the number of equations fixed. As a result, the system of equations becomes more tightly constrained. When the support of Y , i.e., \mathcal{Y} , is finite, the system of equations (4.2)-(4.3) could be solved using a linear programming method, with the drawback that the linear programming approach does not provide intuition about the source of identification power.¹⁵ More importantly, the linear programming approach can no longer be used when Y is continuous, as is the case in our empirical application. To address this issue, we develop a two-step identification approach. The first step provides sharp bounds on the response types. This step involves only the distribution on (D, Z) that has finite support in our framework and can therefore be solved using a linear programming approach, since it does not involve Y which could have continuous support. The second step provides closed-form bounds on the treatment effects of interest, as functions of the sharp bounds on the response types computed in the first step. We show that these two steps provide sharp bounds on our parameters of interest.

4.1. Step 1: Sharp bounds on response type probabilities. In this step, we focus on the partial identification of the distribution of response types. Integrating equations (4.2) and (4.3) over \mathcal{Y} , we obtain the following system of equations, for each d :

$$\mathbb{P}(D = d|Z = 1) = \sum_{d'=0}^K \mathbb{P}[T = (d', d)] \quad (4.4)$$

$$\mathbb{P}(D = d|Z = 0) = \sum_{d'=0}^K \mathbb{P}[T = (d, d')] \quad (4.5)$$

¹⁵If the researcher is interested in analyzing a discrete outcome and wishes to explore this avenue further, she could employ the inferential method developed by Fang et al. (2023).

In general, in the standard IV model, the distribution of response types depends on the full joint distribution of the observed data (Y, D, Z) , not just on the distribution of (D, Z) .¹⁶ This complexity occurs because in the IV framework, the exclusion restriction is imposed, i.e., $Y_{z,d} = Y_{\bullet,d}$. In fact, when this restriction is imposed, the response-type conditional density of $Y_{\bullet,d}$ appears in both equations (4.2) and (4.3), and integrating each equation separately can lead to a loss of information on the response-type probabilities (leading to non-sharp bounds). However, in the absence of the exclusion restriction, each response-type conditional density $f_{Y_{z,d}|T}$ in the system of equations (4.2) and (4.3) only appears in one equation, so the integration step can be performed without losing any information on response-type probabilities. Therefore, we show that in our model, the response-type probabilities are entirely characterized by the distribution of (D, Z) which justifies proceeding in two steps. We will say that a vector \mathbf{v} satisfies (4.4, 4.5) if $(\mathbb{P}[T = (d, d')] : d, d' \in \{0, \dots, K\}) = \mathbf{v}$ is a solution to (4.4, 4.5) for all d .

Lemma 3. *Consider the model (2.5, 2.6). Under Assumption 1, the (sharp) identified set of response-type probabilities is the set of non-negative vectors that satisfy (4.4, 4.5).*

The researcher may also seek to apply additional restrictions on the distribution of response types, including, but not limited to, Assumption 2. For example, one could assume that there are more upward switchers than downward switchers or more stayers than downward switchers. We consider such restrictions as a possible auxiliary assumption. We designate the set of linear constraints that can be applied to the response types as \mathcal{R}_T .

Assumption 3. *[Restriction on response types] Consider that the layers (i.e., firms) are ordered.*

- (i) *[Strong Monotonicity]* $1\{D_1(\omega) = d\} \geq 1\{D_0(\omega) = d'\}$ for $d \geq d'$, or equivalently $\mathbb{P}[T = (d', d)] = 0$ for $d \geq d'$.
- (ii) *[More upward switchers than downward switchers]*
 $\mathbb{P}[T = (d, d')] \geq \mathbb{P}[T = (d', d)]$ for $d \geq d'$.
- (iii) *[More stayers than downward switchers]*
 $\mathbb{P}[T = (d, d)] \geq \mathbb{P}[T = (d', d)]$ for $d \geq d'$.

A notable aspect of Assumptions 2 and 3 is that these restrictions can seamlessly integrate into equations (4.4)-(4.5) as supplementary linear constraints. Consequently, the

¹⁶This has been pointed out by Huber et al. (2017) and is also implicit in the results of Kitagawa (2021). See Theorem 3 in Vayalinkal (2024) for a result characterizing the relationship between outcome distributions and the identified set of response-type probabilities.

process of recovering response types that conform to all these behavioral restrictions simplifies to a feasible linear programming problem. For example, researchers may choose $\mathcal{R}_T = \{\text{Assumption 2}\}$, $\mathcal{R}_T = \{\text{Assumption 2, Assumption 3}(i)\}$, or $\mathcal{R}_T = \{\text{Assumption 2, Assumption 3}\}$.

As mentioned in Lemma 3, equations (4.4, 4.5) sharply characterize the restrictions on the distribution of T imposed by the model (2.5, 2.6). Therefore, the identified set for response type probabilities under model (2.5, 2.6), Assumptions 1, and responses type restrictions \mathcal{R}_T , denoted $\Theta_I(\mathcal{R}_T)$, is simply the set of non-negative vectors that jointly satisfy both \mathcal{R}_T and (4.4, 4.5), i.e.,

$$\Theta_I(\mathcal{R}_T) \equiv \left\{ \mathbf{v} \in \mathbb{R}_{\geq 0}^{(K+1)^2} : (\mathbb{P}[T = (d, d')] : d, d' \in \{0, \dots, K\}) = \mathbf{v} \text{ satisfies } \mathcal{R}_T \text{ and (4.4, 4.5)} \right\}.$$

We also have the following result:

Lemma 4. *Consider the model (2.5, 2.6). Assumption 1 and \mathcal{R}_T are jointly rejected by the data if and only if $\Theta_I(\mathcal{R}_T) = \emptyset$.*

Lemma 4 has significant practical implications. It shows that assessing the validity of model assumptions does not depend on knowledge of the outcome distribution, and therefore, can be reduced to testing the feasibility of a linear program. This property greatly simplifies the implementation of a falsification test for our model.¹⁷ In other words, once we can find a distribution of type that aligns with the assumptions of the model and the observed data on (D, Z) , it is always possible to find a corresponding distribution of potential outcomes $f_{Y_{z,d}|T}(y|d', d)$ that would rationalize the observed joint distribution of (Y, D, Z) .

For simplicity, we introduce the shorthand notation, $p_{d,d'} \equiv \mathbb{P}[T = (d, d')]$, and $\gamma_{d,d'}^z \equiv \frac{p_{d,d'}}{\mathbb{P}(D=d|Z=z)}$. When $\Theta_I(\mathcal{R}_T) \neq \emptyset$, let $\underline{p}_{d,d'}^r$ denote the infimum over all values for $p_{d,d'}$ induced by distributions that belong to $\Theta_I(\mathcal{R}_T)$. Subsequently, we can define: $\underline{\gamma}_{d,d'}^{z,r} = \frac{\underline{p}_{d,d'}^r}{\mathbb{P}(D=d|Z=z)}$. The “ r ” superscript is used to emphasize the dependence of $\underline{p}_{d,d'}^r$ and $\underline{\gamma}_{d,d'}^{z,r}$ on \mathcal{R}_T , the set of assumptions imposed.

For all the potential choices of \mathcal{R}_T considered above, $\Theta_I(\mathcal{R}_T)$ is the set of nonnegative solutions to a linear system. Therefore, $\underline{\gamma}_{d,d'}^{z,r}$ for $d, d' \in \{0, \dots, K\}$ and $z \in \{0, 1\}$ can be obtained as a solution to a linear program. Since the linear system of interest here is generally small, it is also possible to obtain an analytic solution for $\underline{p}_{d,d'}^r$ (and therefore for

¹⁷For example, implementing a (sharp) test of our model is much simpler than the sharp tests for instrument validity and monotonicity proposed by Kitagawa (2015) and Mourifié and Wan (2017).

$\gamma_{d,d'}^{z,r}$) via Fourier-Motzkin elimination. The details of both the computational and analytical approaches are presented in Appendix B.

4.2. Step 2: Sharp bounds on the treatment effects. As evident from equations (4.2) and (4.3), the observed conditional distribution of earnings, $F_{Y|D,Z}(y|d,z)$, can be expressed as a finite mixture of the conditional potential outcome distributions given the response types, $F_{Y_{z,d}|T}(y|l,l')$. More precisely, we have:

$$f_{Y|D=d,Z=1}(y) = \sum_{d'=0}^K \gamma_{d',d}^1 \times f_{Y_{1,d}|T}(y|d',d) \quad (4.6)$$

$$f_{Y|D=d,Z=0}(y) = \sum_{d'=0}^K \gamma_{d,d'}^0 \times f_{Y_{0,d}|T}(y|d,d') \quad (4.7)$$

The unknowns in this mixture are the weights, $\gamma_{d,d'}^z$ for $z \in \{0, 1\}$, and $d, d' \in \{0, \dots, T\}$. In Lee (2009), Assumption 2 implies that the weights are point identified, and establishing identification reduces to recovering the mean average of the mixture components. However, in our scenario, the mixture weights are not point identified under Assumption 2. Even when we strengthen Assumption 2 with Assumption 3, point identification is still not achieved. This underidentification issue arises primarily due to the presence of numerous unobserved types stemming from the multiple layers, i.e., firms. Nevertheless, we can still derive informative bounds on these weights, as elaborated in the previous subsection.

Horowitz and Manski (1995) proposed sharp bounds on the distribution of mixture components in a single-equation mixture model with two components, where the weights are unknown but researchers possess non-trivial bounds for these weights, and Cross and Manski (2002) extended these results to single-equation models with many components. However, the empirical content of our model is characterized by a set of systems of mixture equations, one system for each $d \in \{1, \dots, K\}$, each with up to $(K+1)^2$ components. Importantly, in our setting, the weights are unknown and shared between these systems, introducing a cross-equation dependence that is not present in Horowitz and Manski (1995) or Cross and Manski (2002). We derive the identified set for the weights and then extend the approaches of Horowitz and Manski (1995) and Cross and Manski (2002) to this more general case, deriving closed-form bounds on our key parameters of interest.

Moreover, Lee demonstrated that for continuous outcomes, the bounds proposed by Horowitz and Manski (1995) can be equivalently expressed as the mean of a truncated

distribution. We extend Lee's results by introducing a generalized truncated mean representation that applies regardless of the outcome distribution, whether it is continuous, discrete or mixed.

Hereafter, to simplify our notation and enhance readability, we introduce the following notation. For any d and z , define: $\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\gamma; d, z) \equiv \mathbb{E}[F_{Y|D,Z}^{-1}(U)|U \leq \gamma]$, and $\overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\gamma; d, z) \equiv \mathbb{E}[F_{Y|D,Z}^{-1}(U)|U \geq 1 - \gamma]$, for some $U \sim \text{Uniform}(0, 1)$. We define y_L as the lower bound of the support of Y and y_U as the upper bound.¹⁸

Before stating the main result, we note the following. When the outcome is continuously distributed $\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\gamma; d, z)$ is exactly equal to the truncated mean used in Lee (2009) i.e.,

$$\mathbb{E}[F_{Y|D,Z}^{-1}(U)|U \leq \gamma] = \mathbb{E}[Y|D = d, Z = z, Y \leq F_{Y|D,Z}^{-1}(\gamma)].$$

When the outcome is binary we have

$$\mathbb{E}[F_{Y|D,Z}^{-1}(U)|U \leq \gamma] = \max \left\{ 0, 1 - \frac{1}{\gamma} P[Y = 0|D = d, Z = z] \right\}.$$

This formulation provides a general truncation formula that applies to any type of outcomes, continuous, discrete, or mixed.

Theorem 1. *Suppose that Assumptions 1, and restrictions \mathcal{R}_T hold. Whenever $\Theta_I(\mathcal{R}_T) \neq \emptyset$, then the following bounds are pointwise sharp:*

(i) *Local Controlled Direct Effect (LCDE):*

$$\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d}^{1,r}; d, 1) - \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d}^{0,r}; d, 0) \leq LCDE(d|d, d) \leq \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d}^{1,r}; d, 1) - \underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d}^{0,r}; d, 0),$$

$$\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d',d}^{1,r}; d, 1) - y_U \leq LCDE(d|d', d) \leq \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d',d}^{1,r}; d, 1) - y_L, \text{ for } d \neq d',$$

$$y_L - \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d'}^{0,r}; d, 0) \leq LCDE(d|d', d) \leq y_U - \underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,d'}^{0,r}; d, 0) \text{ for } d \neq d'.$$

(ii) *Local Controlled Indirect Effect (LCIE).*

$$\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{l,d}^{1,r}; d, 1) - y_U \leq LCIE(1, d, d'|l, d) \leq \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{l,d}^{1,r}; d, 1) - y_L, \text{ for } d \neq d' \text{ and any } l,$$

$$\underline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,l}^{0,r}; d, 0) - y_U \leq LCIE(0, d, d'|d, l) \leq \overline{\mathbb{E}}_{F_{Y|D,Z}^{-1}}(\underline{\gamma}_{d,l}^{0,r}; d, 0) - y_L, \text{ for } d \neq d', \text{ and any } l.$$

¹⁸Note that these bounds need not be finite.

(iii) *Aggregate LCDE:*

$$\begin{aligned}
& \inf_{\substack{\{p_{d,d}:d \in \{1, \dots, K\}\} \\ \in \Theta_I(\mathcal{R}_T)}} \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d'=l}^{l'} p_{d',d'}} \left[\mathbb{E}_{F_{Y|D,Z}^{-1}} \left(\frac{p_{d,d}}{\mathbb{P}(D=d|Z=1)}; d, 1 \right) - \bar{\mathbb{E}}_{F_{Y|D,Z}^{-1}} \left(\frac{p_{d,d}}{\mathbb{P}(D=d|Z=0)}; d, 0 \right) \right] \\
& \leq \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d'=l}^{l'} p_{d',d'}} LCDE(d|d, d) \leq \\
& \sup_{\substack{\{p_{d,d}:d \in \{1, \dots, K\}\} \\ \in \Theta_I(\mathcal{R}_T)}} \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d'=l}^{l'} p_{d',d'}} \left[\bar{\mathbb{E}}_{F_{Y|D,Z}^{-1}} \left(\frac{p_{d,d}}{\mathbb{P}(D=d|Z=1)}; d, 1 \right) - \mathbb{E}_{F_{Y|D,Z}^{-1}} \left(\frac{p_{d,d}}{\mathbb{P}(D=d|Z=0)}; d, 0 \right) \right].
\end{aligned}$$

The derivation of the bounds in Theorem 1 comes from extending the Horowitz and Manski (1995) bounding approach summarized in Lemma 5 in Appendix B. However, demonstrating their sharpness presents a considerably more complex challenge. This involves showing that solving equations (4.2) to (4.3) for all $y \in \mathcal{Y}$ and $d \in \{1, \dots, K\}$ while imposing the restrictions defined in \mathcal{R}_T consistently produces the same information as the closed-form bounds presented in Theorem 1. As explained above, the absence of the IV exclusion restrictions facilitates this result.

Theorem 1 (i) indicates that, without additional assumptions on the potential outcome distributions, the derived bounds can best determine the direction (sign) of the within-firm effect at layer d solely for individuals who remain with firm d under any treatment assignment Z , i.e., $\mathbb{E}[Y_{1,d} - Y_{0,d}|T = (d, d)]$. This finding is somewhat intuitive given that these “stayers” are equivalent to the so-called “always-employed” in Lee’s model, where firm heterogeneity is not taken into account. On the other hand, the bounds for those who switch firms due to treatment (“switchers”), such as $\mathbb{E}[Y_{1,d} - Y_{0,d}|T = (d, d')]$ and $\mathbb{E}[Y_{1,d} - Y_{0,d}|T = (d', d)]$ for $d \neq d'$, always include 0. This is because the observed data (Y, D, Z) do not reveal any information on the following unobserved counterfactuals $\mathbb{E}[Y_{0,d}|T = (d', d)]$ and $\mathbb{E}[Y_{1,d}|T = (d, d')]$.

Similarly, Theorem 1 (ii) reveals that in the absence of additional restrictions on the potential outcome distributions, it is impossible to identify the sign of $LCIE(z, d, d'|t) = \mathbb{E}[Y_{z,d} - Y_{z,d'}|T = t]$. This underscores the inherent challenges of identifying some specific treatment effects without imposing further assumptions on the outcome distributions.

Finally, Theorem 1 (iii) presents the closed form bounds that correspond to the weighted average of $LCDE(d|d, d)$, $\sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d'=l}^{l'} p_{d',d}} LCDE(d|d, d)$. These bounds are sharp and take into account the interdependence between equations (4.6) and (4.7).

Remark 2. *To derive bounds on the aggregate LCDE, defined as $\sum_{d=l}^l \frac{p_{d,d}}{\sum_{d'=l}^l p_{d',d'}}$ LCDE($d | d, d$), one might be tempted to adopt a naïve approach by taking a weighted average of the pointwise sharp bounds derived in Theorem 1 (i). However, this approach not only fails to provide sharp bounds on the aggregate quantity but may also yield invalid bounds. We discuss this in more detail in Appendix B.2. The difference between the sharp bounds and naïve “bounds” is illustrated using simulations (Figures 12 and 14) in Appendix A, and using our empirical applications (Figures 5, 8 and 9) in Section 5.*

To better illustrate our theoretical results, in Appendix A, we examine the special case with two firm types, which follows our empirical applications below. We derive closed-form expressions for the bounds on response-type probabilities. Plugging the resulting expression(s) into the bounds given in Theorem 1 (i)-(ii) yields analytic expressions for the bounds on the LCDEs and LCIEs. Through a numerical illustration, we show that our bounds can be sufficiently informative to distinguish cases with a positive within-firm effect from those with no within-firm effect, even when both yield strictly positive Lee bounds.

4.3. Inference. Inference on the causal parameters considered in Theorem 1 can often be performed using existing methods for inference on parameters bounded by truncated conditional expectations. The three sets of bounds in Theorem 1 are known functions of the (D, Z) conditional expectations of Y truncated above or below at particular quantiles. Inference in such settings is complicated by the need to estimate two nuisance parameters: the conditional quantile functions themselves and the truncation quantile level.

In parts (i) and (ii) of Theorem 1, the truncation quantile levels are of the form $\underline{\gamma}_{d,d'}^{z,r} = \frac{p_{d,d'}^r}{\mathbb{P}(D=d|Z=z)}$ and these can often be estimated at a faster rate than the second step (truncated conditional expectation). Moreover, each bound involves either only one truncated conditional expectation, or truncated conditional expectations that can be independently estimated. Therefore, in such cases, inference on the parameters considered in Theorem 1(i)-(ii) can be performed by plugging in the estimators of Lee (2009), Semenova (2020), or Olma (2021) for the truncated conditional expectation(s) and adapting the inference approaches discussed there (subject to the regularity conditions outlined there).¹⁹ These approaches allow for conditioning on, and aggregating over the covariates X , which can lead to much

¹⁹However, $\underline{\gamma}_{d,d'}^{z,r}$ can often be only directionally differentiable with respect to the propensity score vector, as is evident in the closed form expressions for the 2 firms type case provided in Appendix A. In such cases, inference procedures that depend on the bootstrap will need to be adapted to remain valid (see Fang and Santos (2019) for details).

tighter bounds than the unconditional case, as noted by Lee (2009). The approach proposed in Lee (2009) applies when X is finitely supported, whereas the approaches developed by Semenova (2020) and Olma (2021) allow X to be continuous.

Inference on the “aggregate” parameters considered in part (iii) of Theorem 1 is more complicated since the truncation quantile level to be estimated depends on the solution to a higher-dimensional optimization problem involving the outcome distributions. We defer the development of estimation and inference methods for this case to future work.

5. EMPIRICAL APPLICATIONS: JOB CORPS STUDY AND WORKADVANCE RCT

5.1. Application #1: Job Corps Study. Job Corps is the largest residential career training program in the U.S. It is free for participants and targets disadvantaged people aged 16 to 24 with the aim of helping these people become more responsible, employable, and productive citizens (Johnson et al. 1999). Most participants live at a local Job Corps center and complete 440 hours of academic instruction and 700 hours of vocational training. Job Corps also provides job search assistance upon participant completion of the program. The typical participant completes the program over a span of 30 weeks. Job Corps has trained more than two million individuals since its inception under the Economic Opportunity Act of 1964. The program trains over 60,000 enrollees per year, at roughly 130 Job Corps centers nationwide, with an estimated cost of 34,301 USD per enrollee and 57,312 USD per graduate (Liu et al. 2020).

During the mid- to late 1990s, the U.S. Department of Labor funded a randomized evaluation of Job Corps, which was implemented by Mathematica Policy Research, Inc. Existing evaluations of the Job Corps Study include Schochet et al. (2001), Schochet et al. (2008), Lee (2009), and Blanco et al. (2013). The Job Corps Study randomized 80,883 eligible individuals who applied to Job Corps for the first time between November 1994 and December 1995 into two groups: (i) 5,977 individuals into the control group who were embargoed from participating in Job Corps for three years and (ii) 74,906 individuals into the treatment group. Of the 74,906 individuals assigned to treatment, 9,409 individuals were randomly selected for data collection and all control individuals were selected for data collection. The final sample therefore consists of 15,386 participants who were interviewed at the time of random assignment and then subsequently 12, 30, and 48 months after random assignment.

We use the publicly available data from the National Job Corps Study (Schochet et al. 2003). We impose two sample restrictions to address missing values due to interview non-response and sample attrition over time. The first sample restriction, which follows Lee (2009), is to keep individuals who have nonmissing values for weekly earnings and hours for every week following random assignment. Restricting the sample in this manner decreases the sample size to 9,145 individuals (= 3,599 control units + 5,546 treated units).

The second sample restriction is to keep individuals who have non-missing values of health insurance for the weeks of interest (90, 135, 180 and 208). This comes from our classification of firm type based on the provision of health insurance.²⁰ This is motivated by Dey and Flinn (2005) who consider a setting where workers and firms bargain over wages and health insurance and find that, on average, better productivity matches lead to higher wages and the provision of health insurance.²¹ This restriction results in a final sample size of 6,403 individuals (= 2,540 control units + 3,863 treated units).

The three key variables of interest are employment, hourly wage, and provision of health insurance for employed individuals. We follow Lee (2009) by defining employment in a week based on whether an individual has positive earnings in that week and defining the hourly wage in a week by dividing weekly earnings by weekly hours worked. In Appendix C, we present summary statistics for our final sample and demonstrate that, on average, firms that offer these amenities pay higher wages than firms that do not. We also show that the wage distribution for firms that offer amenities stochastically dominates the wage distribution for firms that do not, in both the treatment and control groups. Finally, we show that being randomly assigned to Job Corps led individuals to work at firms with better job amenities compared to the control group. This combined evidence suggests that sample selection is multilayered and motivates the implementation of our sharp bounds to these data.

5.1.1. *Multilayered Bounds for Job Corps Study.* As a first step, we replicate the bounds reported in Lee (2009) which, as discussed above, target the parameter $\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 >$

²⁰Another potentially useful classification of firm type is based on industry since Job Corps targets certain sectors. However, industry is not observed in the public use of the Job Corps data. We consider this type of classification in the next empirical application.

²¹At the time of the National Job Corps study, there were no legal requirements for firms to provide any of these amenities and, conditional on firm provision, federal law generally prohibited discriminatory provision across workers (United States Equal Employment Opportunity Commission 2009). The relevant federal laws at the time of the National Job Corps Study included the following. Title VII of the Civil Rights Act of 1964; the Age Discrimination in Employment Act of 1967; Title I and Title V of the Americans with Disabilities Act of 1990 (United States Equal Employment Opportunity Commission 2009).

$0, D_1 > 0]$.²² In Appendix D.6 we report Lee’s bounds for weeks 90, 135, 180 and 208 along with the trimming proportion $p \equiv \mathbb{P}(AE)$, e.g., the share of the always-employed among individuals receiving job training. Lee focuses on week 90 which produces the tightest bounds $[0.0468, 0.0484]$.²³

We now consider the scenario with two firm types, denoted as L and H . Firms are classified based on the provision of health insurance with H denoting firms that offer health insurance and L denoting firms that do not.²⁴ As in Section 4, we always impose Assumptions 1-2 and then sequentially impose the restrictions in Assumption 3: (i) $p_{H,H} \geq p_{H,L}$ (more stayers than downward switchers) (ii) $\min_t \mathbb{P}[T = t] = p_{H,L}$ (more upward switchers than downward switchers) and (iii) $p_{H,L} = 0$ (strong monotonicity).

Table 1 presents the estimated propensity scores for each week of interest from the National Job Corps Study.²⁵ As expected, in all weeks $\mathbb{P}[D > 0|Z = 1] > \mathbb{P}[D > 0|Z = 0]$ the treated individuals are more likely to be employed. The table also shows that $\mathbb{P}[D = H|D > 0, Z = 1] > \mathbb{P}[D = H|D > 0, Z = 0]$ so that individuals who receive Job Corps training have a higher propensity to be employed in firms that offer health insurance than individuals who do not receive Job Corps training, conditional on employment, consistent with the evidence presented in Appendix Table 11.

TABLE 1. Job Corps: Propensity scores, by week. Health insurance amenity.

	$\mathbb{P}[D = H Z = 0]$	$\mathbb{P}[D = H Z = 1]$	$\mathbb{P}[D = L Z = 0]$	$\mathbb{P}[D = L Z = 1]$
Week 90	0.2239	0.2372	0.2361	0.2229
Week 135	0.2758	0.3037	0.2415	0.2414
Week 180	0.2941	0.3313	0.2462	0.2512
Week 208	0.3142	0.3559	0.2513	0.2509

²²In this section, certain empirical results only presented are for week 90 which follows the preferred specification in Lee (2009). In these cases, tables and figures for other weeks of interest (135, 180 and 208) are presented in Appendix D

²³As we detail in Appendix D.6, these are Lee’s bounds when treating $\ln(\text{hourly wage})$ as a continuous variable, as we do throughout this paper. Lee (2009) uses vintiles of $\ln(\text{hourly wage})$ that produce bounds $[0.0423, 0.0428]$.

²⁴Results for classifying firms based on the provision of pension/retirement benefits and paid vacation are presented in Appendices D.7 and D.8.

²⁵Our sample restriction to keep observations with non-missing amenity values for the weeks of interest drops only employed individuals. This restriction mechanically reduces the propensity scores. To ensure comparability with Lee (2009), we rescale our estimated propensity scores so that the probabilities of employment by treatment status are the same as those reported in Lee (2009).

Using the week 90 propensity scores from Table 1, Figure 4 presents the identified set for $(p_{L,L}, p_{H,H})$. Naturally, incorporating additional restrictions on the response types leads to a tightening of the identified sets. Having characterized the identified set of response-type probabilities, we now present our multilayered bounds.

Recall that under Assumption 2, the support of possible response types is as follows:

$$\mathcal{T} := \{(0, 0), (0, L), (0, H), (L, L), (H, H), (L, H), (H, L)\}.$$

The always-employed (AE) definition used in Lee (2009) therefore combines four different response types: $\{D_0 > 0, D_1 > 0\} = \{(L, L), (H, H), (L, H), (H, L)\} \equiv AE$. We focus on the bounds for stayers, defined as the response types (H, H) and (L, L) .


Figure 5 presents our multilayered bounds for $\mathbb{E}[Y_{1,H} - Y_{0,H} | T = (H, H)]$ and $\mathbb{E}[Y_{1,L} - Y_{0,L} | T = (L, L)]$, along with our aggregate bounds, for weeks 90, 135, 180 and 208. We illustrate the bounds in the baseline case (only Assumptions 1 and 2 is imposed) and also when we sequentially impose the following restrictions: (i) more stayers than downward switchers, (ii) more upward switchers than downward switchers, and (iii) strong monotonicity. As discussed above, the Lee bounds for week 90 are $[0.0468, 0.0484]$. Focusing on the type H firms (firms that offer health insurance), our estimates indicate $\mathbb{E}[Y_{1,H} - Y_{0,H} | T = (H, H)] \in [-2.1415, 2.3907]$. Assuming more stayers than downward switchers tightens these bounds to $[-0.4214, 0.5020]$. Moreover, assuming that (H, L) is the smallest response type, further tightens them to $[-0.0023, 0.0754]$. Finally, assuming strong monotonicity narrows these bounds to $[-0.0018, 0.0750]$.

We find a similar pattern of results for the L -type bounds. These patterns persist across all weeks of interest and also when classifying firms based on the provision of alternative amenities (paid vacation and retirement/pension benefits) as shown in Appendix Figures 34 and 35. Table 2 reports our estimated bounds across weeks.²⁶ The aggregate multilayered bounds are reported in Appendix Table 23.²⁷

Thus, while the conventional Lee bounds are strictly positive, our bounds for the within-firm effects include 0 even under strong assumptions on the response types. This suggests that Lee bounds may capture a pure sorting response to job training rather than a direct wage effect.

²⁶Appendix Tables 24 and 26 provide our estimated within-firm-type bounds when classifying firms based on the provision of paid vacation and retirement/pension benefits, respectively.

²⁷Appendix Tables 25 and 27 provide our estimated aggregate bounds when classifying firms based on the provision of paid vacation and retirement/pension benefits, respectively.



Notes: The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed in addition to Assumption 2. In the two bottom panels of the figure, we shrink the scale of the axes to make the identified set visible.

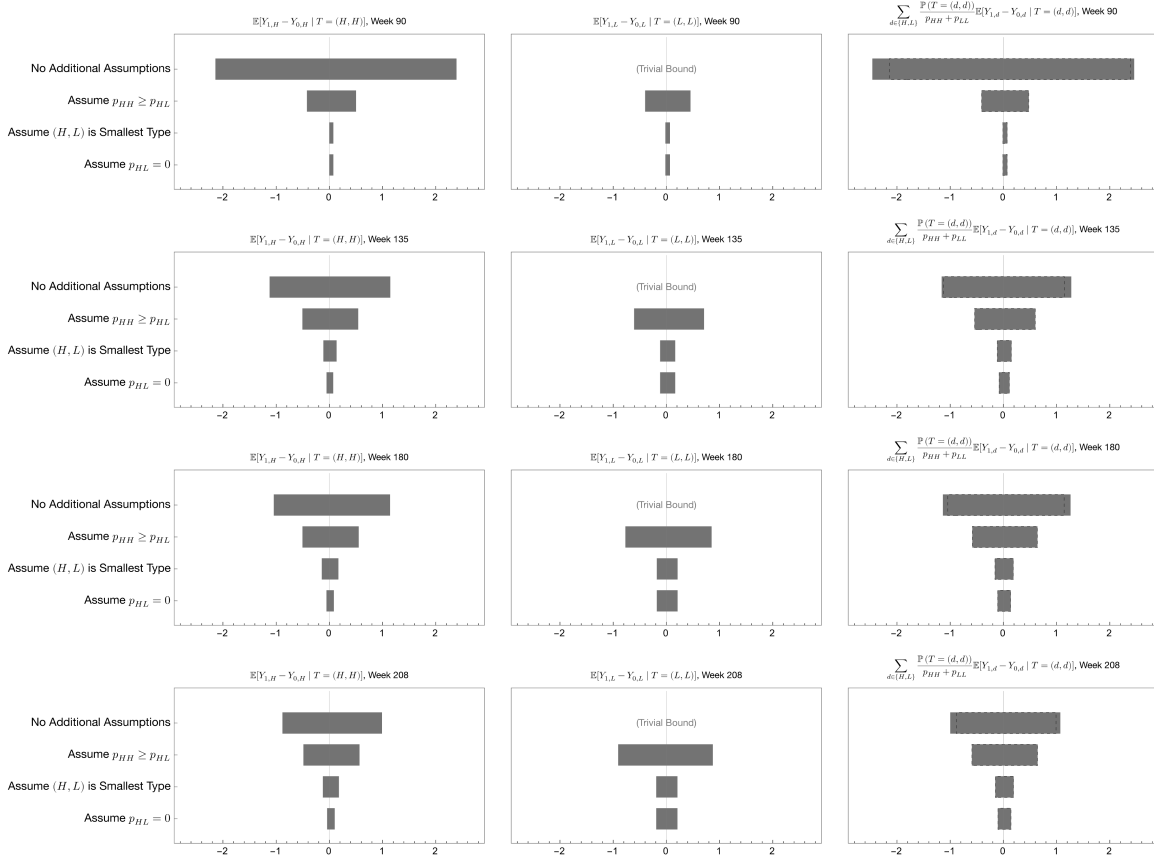
FIGURE 4. Job Corps: Identified set for $(p_{L,L}, p_{H,H})$, week 90. Health insurance amenity.

TABLE 2. Job Corps: Multilayered bounds based on health insurance amenity.

	$p_{H,H}^*$	$p_{L,L}^*$	$\mathbb{E}(Y_{1,H} - Y_{0,H} T = (H, H))$		$\mathbb{E}(Y_{1,L} - Y_{0,L} T = (L, L))$	
			lower	upper	lower	upper
Week 90						
Baseline	0.0010	0.0000	-2.1415	2.3907		
$p_{H,H} \geq p_{H,L}$	0.1120	0.1108	-0.4214	0.5020	-0.4002	0.4542
(H,L) is smallest type	0.2239	0.2228	-0.0023	0.0754	-0.0191	0.0673
$p_{H,L} = 0$	0.2239	0.2228	-0.0018	0.0750	-0.0191	0.0673
Week 135						
Baseline	0.0344	0.0000	-1.1228	1.1454		
$p_{H,H} \geq p_{H,L}$	0.1379	0.0758	-0.5075	0.5433	-0.6064	0.7090
(H,L) is smallest type	0.2619	0.2137	-0.1135	0.1369	-0.1180	0.1672
$p_{H,L} = 0$	0.2758	0.2137	-0.0529	0.0732	-0.1180	0.1672
Week 180						
Baseline	0.0429	0.0000	-1.0454	1.1410		
$p_{H,H} \geq p_{H,L}$	0.1471	0.0620	-0.5063	0.5525	-0.7710	0.8503
(H,L) is smallest type	0.2730	0.2090	-0.1421	0.1704	-0.1806	0.2110
$p_{H,L} = 0$	0.2941	0.2090	-0.0552	0.0854	-0.1806	0.2110
Week 208						
Baseline	0.0633	0.0000	-0.8821	0.9895		
$p_{H,H} \geq p_{H,L}$	0.1571	0.0525	-0.4888	0.5670	-0.9059	0.8730
(H,L) is smallest type	0.2935	0.2096	-0.1217	0.1797	-0.1914	0.2082
$p_{H,L} = 0$	0.3142	0.2096	-0.0430	0.1016	-0.1914	0.2082

Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_t^* is the minimum value of p_t over $\Theta_I(\mathcal{R}_T)$, for the corresponding \mathcal{R}_T .

5.2. Application #2: WorkAdvance RCT. As a second empirical application, we evaluate MDRC's WorkAdvance sectoral employment program. WorkAdvance trains disadvantaged adults with the goal of matching them to high-quality jobs in specific industries with strong labor demand. The MDRC WorkAdvance demonstration featured four different



Notes: Outcome is $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. The first two columns provide the bounds for each firm-level effect and the last column provides bounds on the weighted average of firm-level effects, by week. In all panels, shaded areas indicate sharp bounds; in the last column, thin dashed lines indicate the result of the “naïve approach” of taking the weighted average of firm-level bounds (see Remark 2).

FIGURE 5. Job Corps: Multilayered bounds based on health insurance amenity.

community-based providers in three different locations (New York City, Tulsa, and Northeast Ohio). In this section, we apply our bounds to the Madison Strategies RCT in Tulsa, Oklahoma, and the Towards Employment RCT in Northeast Ohio.^{28,29}

²⁸Our primary dataset is state-level administrative Unemployment Insurance (UI) data, obtained via a confidential data use agreement with MDRC. The administrative UI data for Oklahoma (Madison Strategies RCT) and Ohio (Towards Employment RCT) provide the two-digit North American Industry Classification System (NAICS) code for the industry in which the participant worked, while the New York data (Per Scholas and St. Nicks Alliance RCTs) do not. Therefore, we focus on the Oklahoma and Ohio evaluations.

²⁹For more details on the WorkAdvance RCT and an evaluation of the impacts, see Katz et al. (2022). For Madison Strategies, the reported impact is 12.4 percent on earnings two years after the program. For Towards Employment, the impact is 14 percent.

The Madison Strategies RCT targeted high-quality jobs in Transportation and Manufacturing. The duration of the program ranged from 4 to 32 weeks. The curriculum focused on pre-employment career readiness and sector-specific skills training. It also involved sector-specific placement services, post-employment retention, and advancement services. Upon completion of the program, participants received the appropriate certification.

The Madison Strategies RCT enrollment period was from June 2011 to June 2013. It targeted low-income adults who met the following skill requirements: (i) tested at the eighth grade level; (ii) passed a behavioral assessment; (iii) passed mechanical aptitude and manual dexterity exams; (iv) had a driver's license. The control group was not eligible to receive WorkAdvance services. Eligible potential participants (697 units) were randomly assigned to treatment (353 units) and control (344 units).

The Towards Employment RCT targeted jobs in Health Care and Manufacturing. The target participants were low-income adults who: (i) tested at the sixth to tenth grade level; (ii) passed background and drug tests. The curriculum, program duration, and enrollment period were the same as Madison Strategies. Eligible potential participants (698 units) were randomly assigned to treatment (349 units) and control (349 units).

The key variables for our analysis of both RCTs are quarterly wages along with employment status and sector. We define quarterly wages as quarterly earnings subject to UI and follow both Lee (2009)'s and our evaluation of the Job Corps RCT in defining quarterly employment status based on whether an individual has positive earnings in a quarter.³⁰ We assume that $d = H$ if the firm of employment is in the target sector and $d = L$ if the firm is not. All of our results focus on 8 quarters post-random assignment.

Appendix Tables 28 and 29 show summary statistics at baseline, as well earnings and employment outcomes 8 quarters post-randomization, for Madison Strategies and Towards Employment, respectively. For Madison Strategies, the employment rates 8 quarters post-randomization are quite similar between the treatment group (0.67) and the control group (0.66). However, the composition of employment by target sector differs meaningfully by treatment status: among the treatment group, 44 percent are employed in the target sector whereas, among the control group, only 31 percent are. Towards Employment increased overall employment from 0.62 to 0.69 and the share of employment in the target sector from 0.47 to 0.51. Table 3 shows the propensity score estimates for both Madison Strategies and Towards Employment which further illustrate the sorting effects of both RCTs. As was the

³⁰To interpret our estimates as reflecting wage effects, this requires that job training does not affect hours of work.

case for the Job Corps study, this combined evidence suggests multilayered sample selection and again motivates the use of our bounds.

Appendix Figure 36 shows that average wages are higher in the target sector and Appendix Figure 37 (Madison Strategies) and Appendix Figure 38 (Towards Employment) show that the wage distributions in the target sector stochastically dominate the wage distributions in non-targeted sectors, conditional on treatment status.

TABLE 3. Propensity scores for WorkAdvance RCTs


	$\mathbb{P}[D = H Z = 0]$	$\mathbb{P}[D = H Z = 1]$	$\mathbb{P}[D = L Z = 0]$	$\mathbb{P}[D = L Z = 1]$
Madison Strategies	0.2035	0.2975	0.4535	0.3711
Towards Employment	0.2923	0.3524	0.3238	0.3410

Figure 6 and Figure 7 show the identified sets for $(p_{L,L}, p_{H,H})$ for Madison Strategies and Towards Employment, respectively. Although sorting appears to be stronger in the WorkAdvance RCT than in the Job Corps study, $p_{H,H}^*$ and $p_{L,L}^*$ are fairly similar. This is driven by differences in $p_{0,0}$, which is a point identified by $P(D = 0|Z = 1)$. In Job Corps, $p_{0,0} = 0.54$. By comparison, $p_{0,0} = 0.33$ in Madison Strategies and $p_{0,0} = 0.31$ in Towards Employment. Thus, the reason for the similarities in $p_{H,H}^*$ and $p_{L,L}^*$ is driven by much stronger sorting (larger p_{LH}) in WorkAdvance. All else being equal, this attenuates $p_{d,d}^*$ and increases the bounds.

TABLE 4. Multilayered bounds for WorkAdvance RCT.

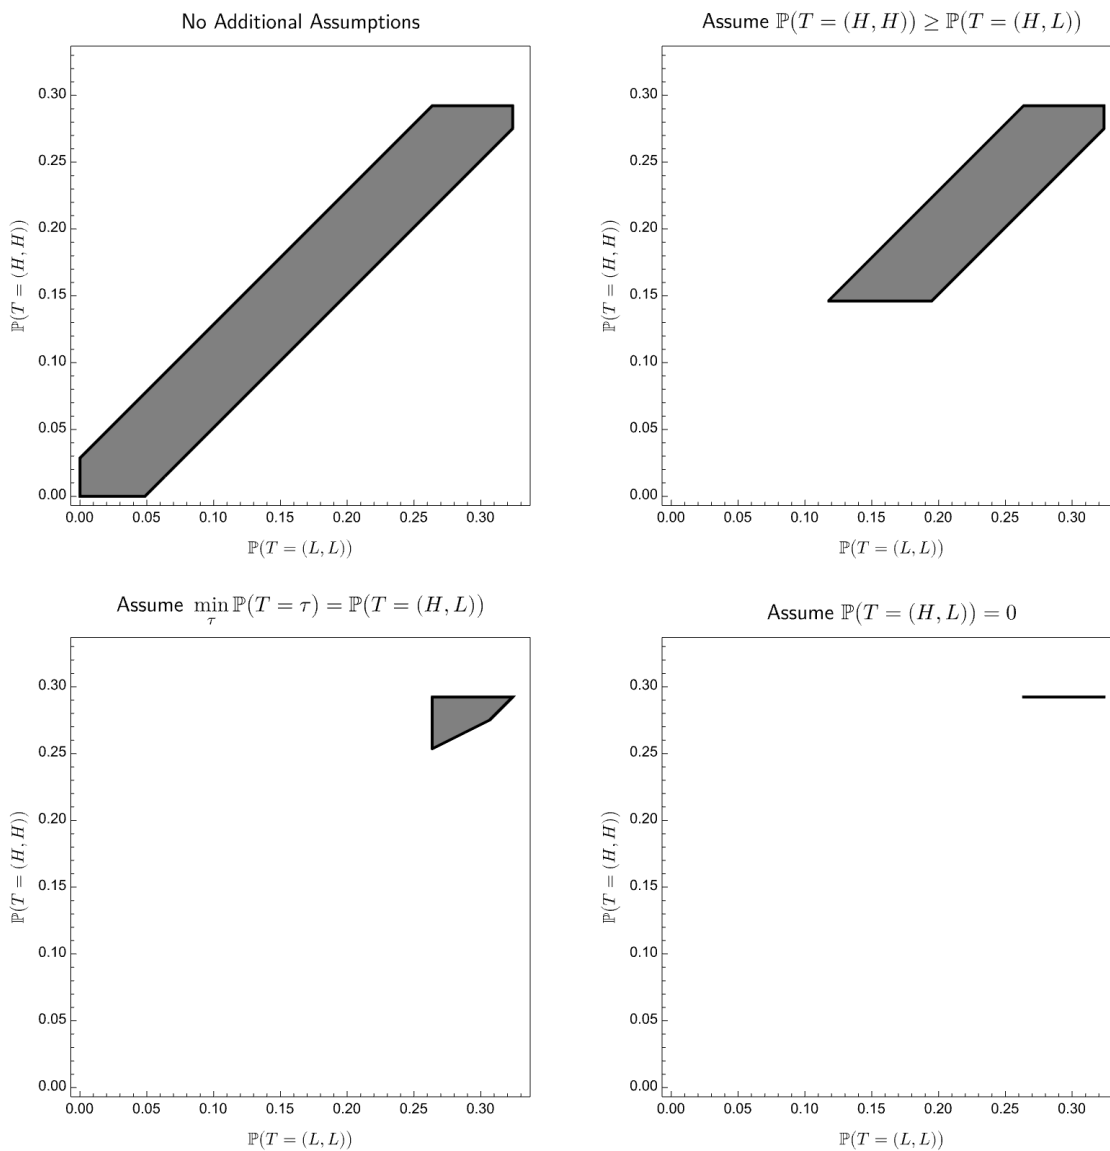
	$p_{H,H}^*$	$p_{L,L}^*$	$\mathbb{E}(Y_{1,H} - Y_{0,H} T = (H, H))$		$\mathbb{E}(Y_{1,L} - Y_{0,L} T = (L, L))$	
			lower	upper	lower	upper
Madison Strategies						
Baseline	0.0000	0.1560			-6172.4634	7139.9806
$p_{H,H} \geq p_{H,L}$	0.1017	0.2578	-7380.1460	5628.8654	-3353.0004	4562.1007
(H,L) is smallest type	0.1977	0.3595	-2652.6847	1216.5151	-798.7005	2016.2079
$p_{H,L} = 0$	0.2035	0.3595	-2353.4366	806.8709	-798.7005	2016.2079
Towards Employment						
Baseline	0.0000	0.0000				
$p_{H,H} \geq p_{H,L}$	0.1461	0.1175	-5784.7451	6488.1176	-6510.9268	5975.5610
(H,L) is smallest type	0.2536	0.2636	-2044.9605	3064.8644	-2335.5870	1932.3913
$p_{H,L} = 0$	0.2923	0.2636	-740.4216	1356.9608	-2335.5870	1932.3913

Notes: Outcome is quarterly wages. p_t^* is the minimum value of p_t over the identified set for response-types under the given assumption.



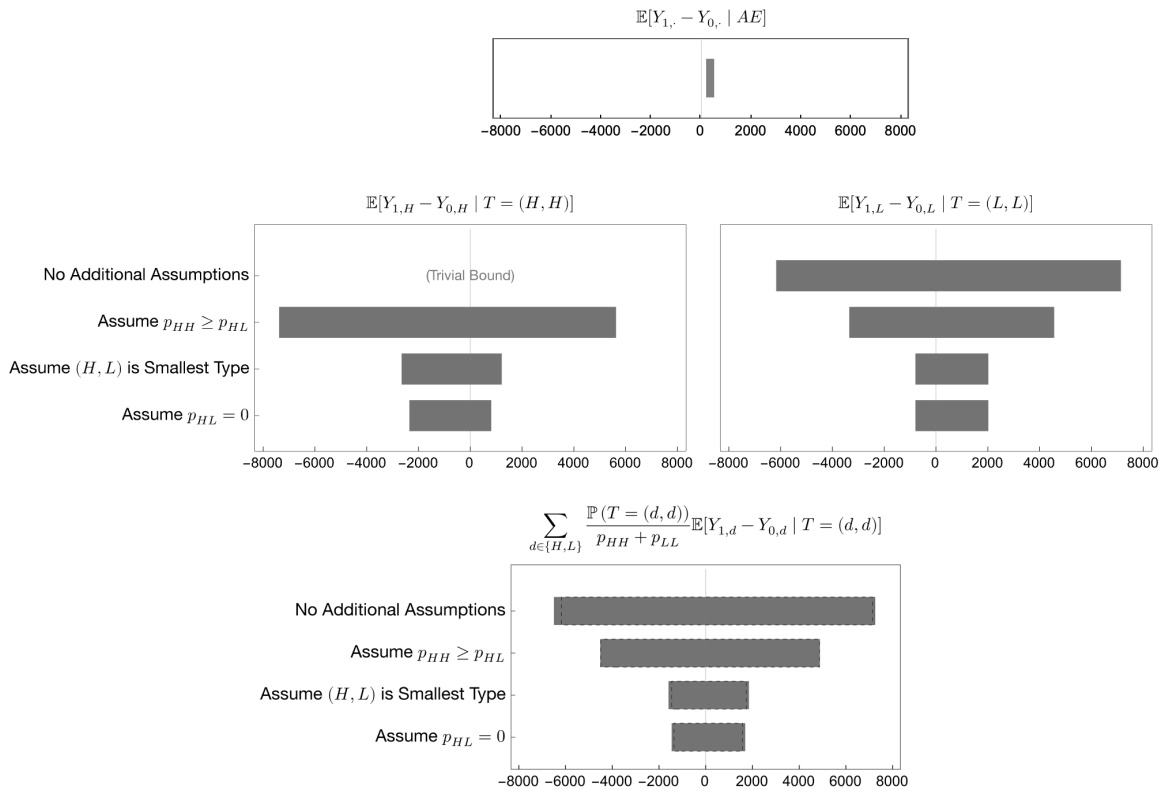
Notes: The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed in addition to Assumption 2.

FIGURE 6. Madison Strategies RCT: Identified set for $(p_{L,L}, p_{H,H})$.



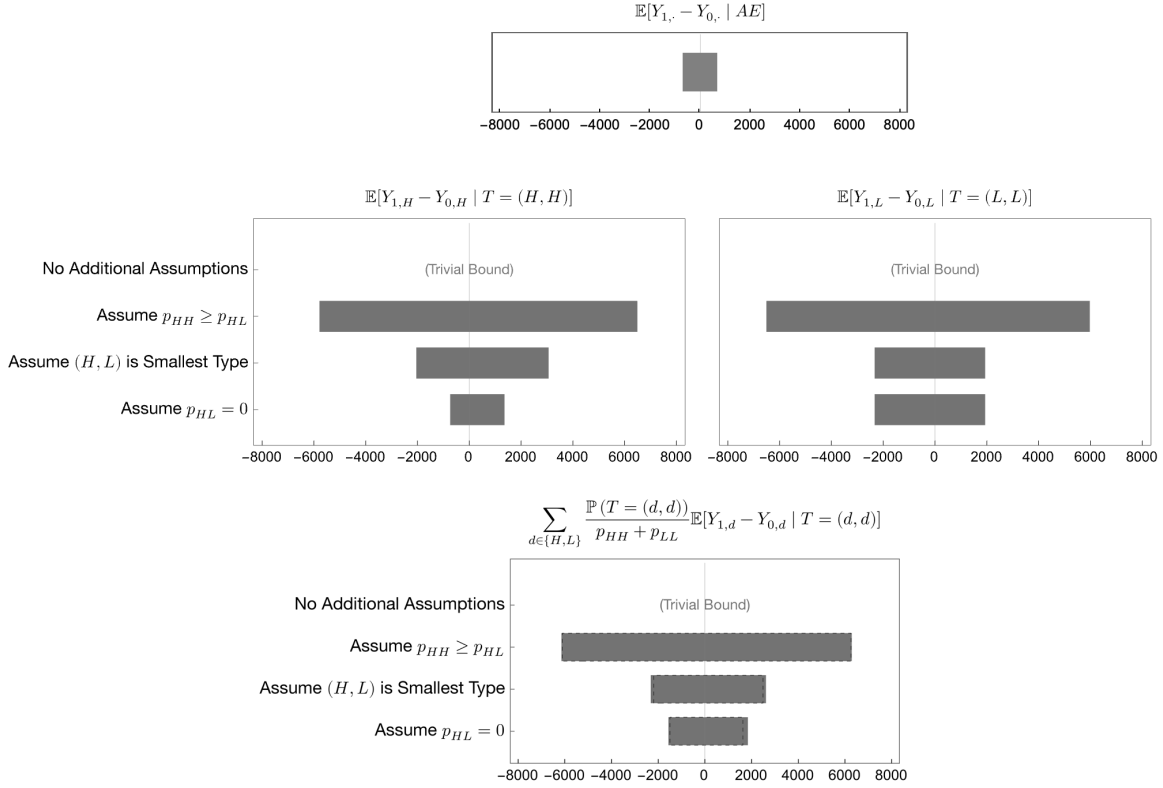
Notes: The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed in addition to Assumption 2.

FIGURE 7. Towards Employment RCT: Identified set for $(p_{L,L}, p_{H,H})$.



Notes: Top panel illustrates Lee (2009) bounds, under Assumptions 1 and 2. Remaining panels illustrate the multilayered bounds under various sets of additional assumptions: the middle panels provide the bounds for each firm-level effect, and the bottom panel provides bounds on the weighted average of firm-level effects. In all panels, shaded areas indicate sharp bounds; in the bottom panel, thin dashed lines indicate the result of the “naïve approach” of taking the weighted average of firm-level bounds (see Remark 2).

FIGURE 8. Bounds for the Madison Strategies RCT.



Notes: Top panel illustrates Lee (2009) bounds, under Assumptions 1 and 2. Remaining panels illustrate the multilayered bounds under various sets of additional assumptions: the middle panels provide the bounds for each firm-level effect, and the bottom panel provides bounds on the weighted average of firm-level effects. In all panels, shaded areas indicate sharp bounds; in the bottom panel, thin dashed lines indicate the result of the “naïve approach” of taking the weighted average of firm-level bounds (see Remark 2).

FIGURE 9. Bounds for the Towards Employment RCT.

Figures 8 and 9 present Lee bounds, our multilayered bounds for $\mathbb{E}[Y_{1,H} - Y_{0,H} | T = (H, H)]$ and $\mathbb{E}[Y_{1,L} - Y_{0,L} | T = (L, L)]$, along with our aggregate bounds, for Madison Strategies and Towards Employment, respectively. As before, we present our multilayered bounds in the baseline case (only Assumptions 1 and 2 are imposed) and then sequentially impose the following restrictions: (i) more stayers than downward switchers, (ii) more upward switchers than downward switchers, and (iii) strong monotonicity. Even under the most restrictive assumption of strong monotonicity, the within-firm bounds include 0 for both sites, suggesting that the impact of WorkAdvance on wages is primarily driven by sorting to the target sector. Notably, this contrasts with Lee bounds in the case of

Madison Strategies, which are positive and tight. Table 4 presents the within-firm-type multilayered bounds for Madison Strategies (top panel) and Towards Employment (bottom panel). Appendix Tables 30 and 31 present the aggregate multilayered bounds and Lee bounds, respectively, for both sites.

6. CONCLUSION

This paper develops a new methodology to partially identify the causal effect of job training on wages in the presence of multilayered sample selection. We define new treatment effects that operate within and between firms and provide a new identification approach that extends the Horowitz and Manski (1995) bounds. As a proof of concept, we show how to empirically implement these bounds by considering applications to the Job Corps Study and the WorkAdvance RCT.

Although we consider our approach in the context of job training where a layer corresponds to a firm, we view it as naturally extending to other settings. In particular, it applies to any setting where sample selection is multilayered. As an example, consider a setting where a researcher is interested in estimating the causal effect of a tuition subsidy on labor market outcomes.³¹ The subsidy may have an effect on the type of institution that an individual enrolls in and graduates from. If earnings depend on institutional quality, part of the earnings effects of the subsidy could reflect the value-added of institutions that are affected by the subsidy.

Although our framework has focused mainly on nonparametric (partial) identification, we are currently reexamining the classic parametric sample selection approach of Heckman (1979) in the context of multilayered sample selection, as well as its semiparametric version discussed in Honoré and Hu (2020). By imposing additional structure on the unobservables, this approach has the potential to significantly tighten the bounds and may achieve point identification of causal effects.

³¹Bettinger et al. (2019) evaluate the impact of California’s state-based financial aid on long-run earnings.

REFERENCES

- ABOWD, JOHN M., FRANCIS KRAMARZ, AND DAVID N. MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67, 251–333.
- AHN, HYUNGTAIK AND JAMES L. POWELL (1993): “Semiparametric estimation of censored selection models with a nonparametric selection mechanism,” *Journal of Econometrics*, 58, 3–29.
- ANDERSSON, FREDRIK, HARRY J. HOLZER, JULIA I. LANE, DAVID ROSENBLUM, AND JEFFREY SMITH (2022): “Does Federally-Funded Job Training Work? Nonexperimental Estimates of WIA Training Impacts Using Longitudinal Data on Workers and Firms,” *Journal of Human Resources*, 0816–8185R1.
- BETTINGER, ERIC, ODED GURANTZ, LAURA KAWANO, BRUCE SACERDOTE, AND MICHAEL STEVENS (2019): “The Long-Run Impacts of Financial Aid: Evidence from California’s Cal Grant,” *American Economic Journal: Economic Policy*, 11, 64–94.
- BIANCHI, NICOLA AND MICHELA GIORCELLI (2022): “The Dynamics and Spillovers of Management Interventions: Evidence from the Training within Industry Program,” *Journal of Political Economy*, 130, 1630–1675.
- BLANCO, GERMAN, CARLOS A FLORES, AND ALFONSO FLORES-LAGUNES (2013): “The Effects of Job Corps Training on Wages of Adolescents and Young Adults,” *American Economic Review*, 103, 418–422.
- BONHOMME, STÉPHANE, KERSTIN HOLZHEU, THIBAUT LAMADON, ELENA MANRESA, MAGNE MOGSTAD, AND BRADLEY SETZLER (2023): “How Much Should We Trust Estimates of Firm Effects and Worker Sorting?” *Journal of Labor Economics*, 41, 291–322.
- BONHOMME, STÉPHANE, THIBAUT LAMADON, AND ELENA MANRESA (2019): “A Distributional Framework for Matched Employer Employee Data,” *Econometrica*, 87, 699–739.
- CANAY, IVAN, ANDRES SANTOS, AND AZEEM M SHAIKH (2013): “On the Testability of Identification in Some Nonparametric Models With Endogeneity,” *Econometrica*, 81, 2535–2559.
- CARD, DAVID, ANA RUTE CARDOSO, AND PATRICK KLINE (2016): “Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women*,” *The Quarterly Journal of Economics*, 131, 633–686.
- CARD, DAVID, JÖRG HEINING, AND PATRICK KLINE (2013): “Workplace Heterogeneity and the Rise of West German Wage Inequality*,” *The Quarterly Journal of Economics*, 128, 967–1015.

- CARD, DAVID, JOCHEN KLUVE, AND ANDREA WEBER (2010): “Active Labour Market Policy Evaluations: A Meta-Analysis,” *The Economic Journal*, 120, F452–F477.
- (2018): “What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations,” *Journal of the European Economic Association*, 16, 894–931.
- CROSS, PHILIP J. AND CHARLES F. MANSKI (2002): “Regressions, Short and Long,” *Econometrica*, 70, 357–368.
- CULLEN, ZOË B. AND BOBAK PAKZAD-HURSON (2023): “Equilibrium Effects of Pay Transparency,” *Econometrica*, 91, 765–802.
- DARUICH, DIEGO, SABRINA DI ADDARIO, AND RAFFAELE SAGGIO (2023): “The Effects of Partial Employment Protection Reforms: Evidence from Italy,” *Review of Economic Studies*, 90, 2880–2942.
- DEY, MATTHEW S. AND CHRISTOPHER J. FLINN (2005): “An Equilibrium Model of Health Insurance Provision and Wage Determination,” *Econometrica*, 73, 571–627.
- D’HAULTFOEUILLE, XAVIER (2011): “On the Completeness Condition in Nonparametric Instrumental Problems,” *Econometric Theory*, 27, 460–471.
- FANG, ZHENG AND ANDRES SANTOS (2019): “Inference on Directionally Differentiable Functions,” *The Review of Economic Studies*, 86, 377–412.
- FANG, ZHENG, ANDRES SANTOS, AZEEM M. SHAIKH, AND ALEXANDER TORGOVITSKY (2023): “Inference for Large-Scale Linear Systems With Known Coefficients,” *Econometrica*, 91, 299–327.
- FINK, GÜNTHER, B. KELSEY JACK, AND FELIX MASIYE (2020): “Seasonal Liquidity, Rural Labor Markets, and Agricultural Production,” *American Economic Review*, 110, 3351–3392.
- FISMAN, RAYMOND, DANIEL PARAVISINI, AND VIKRANT VIG (2017): “Cultural Proximity and Loan Outcomes,” *American Economic Review*, 107, 457–492.
- GALLANT, A. RONALD AND DOUGLAS W. NYCHKA (1987): “Semi-Nonparametric Maximum Likelihood Estimation,” *Econometrica*, 55, 363.
- GERARD, FRANÇOIS, LORENZO LAGOS, EDSON SEVERNINI, AND DAVID CARD (2021): “Assortative Matching or Exclusionary Hiring? The Impact of Employment and Pay Policies on Racial Wage Differences in Brazil,” *American Economic Review*, 111, 3418–3457.
- GIORCELLI, MICHELA (2019): “The Long-Term Effects of Management and Technology Transfers,” *American Economic Review*, 109, 121–152.
- GOTTSCHALK, PETER, DAVID GREEN, AND BEN SAND (2014): “Taking Selection to Task: Bounds on Trends in Occupational Task Prices for the U.S., 1984-2014,” *Unpublished*

Manuscript.

- HECKMAN, JAMES J. (1979): “Sample Selection Bias as a Specification Error,” *Econometrica*, 47, 153.
- HECKMAN, JAMES J., ROBERT J. LALONDE, AND JEFFREY A. SMITH (1999): “The Economics and Econometrics of Active Labor Market Programs,” in *Handbook of Labor Economics*, Elsevier, vol. 3, 1865–2097.
- HECKMAN, JAMES J. AND RODRIGO PINTO (2018): “Unordered Monotonicity,” *Econometrica*, 86, 1–35.
- HENDREN, NATHANIEL AND BEN SPRUNG-KEYSER (2020): “A Unified Welfare Analysis of Government Policies*,” *The Quarterly Journal of Economics*, 135, 1209–1318.
- HONORÉ, BO E. AND LUOJIA HU (2020): “Selection Without Exclusion,” *Econometrica*, 88, 1007–1029.
- HOROWITZ, JOEL L. AND CHARLES F. MANSKI (1995): “Identification and Robustness with Contaminated and Corrupted Data,” *Econometrica*, 63, 281.
- HUBER, MARTIN, LUKAS LAFFERS, AND GIOVANNI MELLACE (2017): “Sharp IV Bounds on Average Treatment Effects on the Treated and Other Populations Under Endogeneity and Noncompliance,” *Journal of Applied Econometrics*, 32, 56–79.
- IMBENS, GUIDO W. AND JOSHUA D. ANGRIST (1994): “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62, 467.
- JOHNSON, TERRY, MARK GRITZ, RUSSELL JACKSON, JOHN BURGHARDT, CAROL BOUSSY, JAN LEONARD, AND CARLYN ORIAN (1999): “National Job Corps Study: Report on the Process Analysis,” Tech. rep., Mathematica Policy Research Inc., Princeton, NJ.
- KATZ, LAWRENCE F., JONATHAN ROTH, RICHARD HENDRA, AND KELSEY SCHABERG (2022): “Why Do Sectoral Employment Programs Work? Lessons from WorkAdvance,” *Journal of Labor Economics*, 40, S249–S291.
- KITAGAWA, TORU (2015): “A Test for Instrument Validity,” *Econometrica*, 83, 2043–2063.
- (2021): “The identification region of the potential outcome distributions under instrument independence,” *Journal of Econometrics*, 225, 231–253.
- KWON, SOONWOO AND JONATHAN ROTH (2024): “Testing Mechanisms,” *arXiv preprint arXiv:2404.11739*.
- LACHOWSKA, MARTA, ALEXANDRE MAS, AND STEPHEN A. WOODBURY (2020): “Sources of Displaced Workers’ Long-Term Earnings Losses,” *American Economic Review*, 110, 3231–3266.

- LAMADON, THIBAUT, MAGNE MOGSTAD, AND BRADLEY SETZLER (2022): “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, 112, 169–212.
- LEE, DAVID S. (2009): “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects,” *Review of Economic Studies*, 76, 1071–1102.
- LI, QI AND JEFFREY RACINE (2007): *Nonparametric econometrics: theory and practice*, Princeton Oxford: Princeton University Press.
- LIU, ALBERT, CLAYTON PERRY, ROBBIE SKINNER, AND SCOTT CODY (2020): “Estimating Job Corps Cost per Enrollee and Cost per Graduate,” Tech. rep., Insight Policy Research, Arlington, VA.
- MOURIFIÉ, ISMAEL AND YUANYUAN WAN (2017): “Testing Local Average Treatment Effect Assumptions,” *The Review of Economics and Statistics*, 99, 305–313.
- NEWHEY, WHITNEY, JAMES POWELL, AND WALKER (1990): “Semiparametric Estimation of Selection Models: Some Empirical Results,” *American Economic Review*, 80, 324–328.
- OLMA, TOMASZ (2021): “Nonparametric Estimation of Truncated Conditional Expectation Functions,” Version Number: 1.
- PEARL, JUDEA (2001): “Direct and Indirect Effects,” Tech. rep.
- (2009): “Causal inference in statistics: An overview,” *Statistics Surveys*, 3.
- ROBINS, JAMES M. AND SANDER GREENLAND (1992): “Identifiability and Exchangeability for Direct and Indirect Effects:,” *Epidemiology*, 3, 143–155.
- SCHMIEDER, JOHANNES F., TILL VON WACHTER, AND JÖRG HEINING (2023): “The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany,” *American Economic Review*, 113, 1208–1254.
- SCHOCHET, PETER, JEANNE BELLOTTI, RUO-JIAO CAO, STEVEN GLAZERMAN, APRIL GRADY, MARK GRITZ, SHEENA MCCONNELL, TERRY JOHNSON, AND JOHN BURGHARDT (2003): “National Job Corps Study: Data Documentation and Public Use Files: Volume I,” .
- SCHOCHET, PETER, JOHN BURGHARDT, AND STEVEN GLAZERMAN (2001): “National Job Corps Study: The Impacts of Job Corps on Participants’ Employment and Related Outcomes,” Tech. rep., Mathematica Policy Research Inc., Princeton, NJ.
- SCHOCHET, PETER Z, JOHN BURGHARDT, AND SHEENA MCCONNELL (2008): “Does Job Corps Work? Impact Findings from the National Job Corps Study,” *American Economic Review*, 98, 1864–1886.
- SEMENOVA, VIRIA (2020): “Generalized Lee Bounds,” Publisher: [object Object] Version Number: 3.

- SONG, JAE, DAVID J PRICE, FATIH GUVENEN, NICHOLAS BLOOM, AND TILL VON WACHTER (2019): “Firming Up Inequality*,” *The Quarterly Journal of Economics*, 134, 1–50.
- SŁOCZYŃSKI, TYMON (2020): “When Should We (Not) Interpret Linear IV Estimands as LATE?” Publisher: [object Object] Version Number: 6.
- THE COUNCIL OF ECONOMIC ADVISERS (2019): “Government Employment and Training Programs: Assessing the Evidence on their Performance,” Tech. rep.
- UNITED STATES EQUAL EMPLOYMENT OPPORTUNITY COMMISSION (2009): “Federal Laws Prohibiting Job Discrimination Questions and Answers,” Tech. rep., United States Equal Employment Opportunity Commission, Washington, DC.
- VAYALINKAL, ATOM (2024): “Sharp Identification Regions in General Selection Models with (Un)ordered Treatments and Discrete Instruments,” *Unpublished Manuscript, University of Toronto*.
- ZUO, SHUOZHI, DEBASHIS GHOSH, PENG DING, AND FAN YANG (2022): “Mediation analysis with the mediator and outcome missing not at random,” Version Number: 3.

APPENDIX A. 2 FIRM TYPES CASE: A NUMERICAL ILLUSTRATION

To provide intuition for our bounding approach, in this section we consider a scenario involving two types of firms. We assume that the firms are identical within each type. Firms are categorized as high type (H) or low type (L). Our objective is to investigate whether the instrument has a causal effect on wages within each firm type. Under Assumption 2, the support of possible response types is:

$$\mathcal{T} := \{(0, 0), (0, L), (0, H), (L, L), (H, H), (L, H), (H, L)\}.$$

Thus, the always-employed (AE) encompasses four distinct response types:

$$\{D_0 > 0, D_1 > 0\} = \{(L, L), (H, H), (L, H), (H, L)\} \equiv AE.$$

Lee bounds applied to this setting correspond to:

$$\begin{aligned} \theta^\ell &\leq \frac{p_{L,L} + p_{H,L}}{\mathbb{P}(AE)} \mathbb{E}[Y_{1,L} - Y_{0,L} | T \in \{(L, L), (H, L)\}] \\ &+ \frac{p_{H,H} + p_{L,H}}{\mathbb{P}(AE)} \mathbb{E}[Y_{1,H} - Y_{0,H} | T \in \{(H, H), (L, H)\}] + \frac{p_{L,H}}{\mathbb{P}(AE)} \mathbb{E}[Y_{0,H} - Y_{0,L} | T = (L, H)] + \\ &\quad \frac{p_{H,L}}{\mathbb{P}(AE)} \mathbb{E}[Y_{0,L} - Y_{0,H} | T = (H, L)] \leq \bar{\theta}^\ell. \end{aligned}$$

To establish bounds on our causal effects of interest, we first characterize identification of the response-type probabilities: $\{p_t : t \in \mathcal{T}\}$. Using information on (D, Z) only, $\{p_t : t \in \mathcal{T}\}$ has to satisfy equations (4.4, 4.5) from step 1 above. In addition, if we impose Assumption 2, i.e., $\mathcal{R}_T = \{\text{Assumption 2}\}$, we can show that the identified set for the response types in this

simple case is characterized by non-negative solutions to the following set of (in)equalities:

$$p_{0,0} = 1 - \mathbb{P}(D = H | Z = 1) - \mathbb{P}(D = L | Z = 1), \quad (\text{A.1})$$

$$p_{0,L} = \mathbb{P}(D = L | Z = 1) - \mathbb{P}(D = H | Z = 0) + p_{H,H} - p_{L,L}, \quad (\text{A.2})$$

$$p_{0,H} = \mathbb{P}(D = H | Z = 1) - \mathbb{P}(D = L | Z = 0) + p_{L,L} - p_{H,H}, \quad (\text{A.3})$$

$$p_{L,H} = \mathbb{P}(D = L | Z = 0) - p_{L,L}, \quad (\text{A.4})$$

$$p_{H,L} = \mathbb{P}(D = H | Z = 0) - p_{H,H}, \quad (\text{A.5})$$

$$\begin{aligned} \max\{0, \mathbb{P}(D = H | Z = 0) - \mathbb{P}(D = L | Z = 1)\} \leq \\ p_{H,H} \leq \min\{\mathbb{P}(D = H | Z = 0), \mathbb{P}(D = H, Z = 1)\} \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} \max\{0, \mathbb{P}(D = L | Z = 0) - \mathbb{P}(D = H | Z = 1)\} \leq \\ p_{L,L} \leq \min\{\mathbb{P}(D = L | Z = 0), \mathbb{P}(D = L, Z = 1)\}. \end{aligned} \quad (\text{A.7})$$

More precisely, we can show that

$$\Theta_I(\mathcal{R}_T) = \{(p_t : t \in \mathcal{T}) \geq 0 : \text{ such that eqs (A.1) to (A.7) are satisfied}\}.$$

In this case, it is also straightforward to show that $\Theta_I(\mathcal{R}_T) \neq \emptyset$ if and only if the propensity scores satisfy

$$\mathbb{P}(D = 0 | Z = 1) \leq \mathbb{P}(D = 0 | Z = 0). \quad (\text{A.8})$$

Having defined the identified set for response types, we proceed to construct bounds on our treatment effects of interest. In this particular case, we have:

$$\begin{aligned} \underline{\gamma}_{H,H}^z &= \frac{\max\{0, \mathbb{P}(D = H | Z = 0) - \mathbb{P}(D = L | Z = 1)\}}{\mathbb{P}(D = H | Z = z)}, \text{ for } z \in \{0, 1\}, \\ \underline{\gamma}_{L,L}^z &= \frac{\max\{0, \mathbb{P}(D = L | Z = 0) - \mathbb{P}(D = H | Z = 1)\}}{\mathbb{P}(D = L | Z = z)}, \text{ for } z \in \{0, 1\}, \end{aligned}$$

One can then apply the closed-form formula from Theorem 1 to establish bounds on the treatment effects of interest. A unique aspect of our methodology is that imposing additional restrictions on response types does not alter the form of the bounds in Theorem 1. The bounds remain valid (and sharp) for the new set of restrictions, as long as $\underline{\gamma}_{d,d}^z$ is adjusted according to the new restrictions.³²

³²That is, if $\underline{\gamma}_{d,d}^z$ is updated to be the result of optimizing over the new $\Theta_I(\mathcal{R}_T)$.

In the numerical illustration below, we describe $\Theta_I(\mathcal{R}_T)$ and demonstrate how imposing further assumptions on response types can significantly refine $\Theta_I(\mathcal{R}_T)$. More precisely, we consider the following additional assumptions.

- (i) $p_{H,H} \geq p_{H,L}$: remaining within a high-type firm is more probable than transitioning from a high-type to a low-type firm as a consequence of the treatment.
- (ii) $\min_t p_t = p_{H,L}$: the smallest proportion of response type consists of individuals moving from a high-type to a low-type firm due to the treatment.
- (iii) $p_{H,L} = 0$: the absence of transitions from high-type to low-type firms as a result of treatment.

In this case, as shown by (A.1)-(A.5), the response-type probability p_t , for each $t \in \{(0,0), (0,L), (0,H), (L,H), (H,L)\}$, can be represented as a linear function of $p_{H,H}$ and $p_{L,L}$, which are themselves only set identified; in other words, $\Theta_I(\mathcal{R}_T)$ is parameterized by $(p_{H,H}, p_{L,L})$. Hence, our forthcoming discussion will focus primarily on illustrating the projection of $\Theta_I(\mathcal{R}_T)$ with respect to the coordinates of $(p_{H,H}, p_{L,L})$.

A.0.1. *Data Generating Process.* We first generate propensity scores to be consistent with the data in Lee (2009). This is reported in Table 5.

TABLE 5. Propensity scores for numerical illustrations

Job Training	$P(D = L Z = z)$	$P(D = H Z = z)$
$Z = 1$	0.302886	0.408114
$Z = 0$	0.313959	0.373041

The data are generated such that the true values for $p_{H,H}$, and $p_{L,L}$ are:

$$\begin{aligned} p_{H,H} &= \mathbb{P}(D = H | Z = 0) = 0.373041, \\ p_{L,L} &= 0.278886. \end{aligned}$$

Next, we randomly generate the outcomes.³³ For each type $t \in \mathcal{T}$, denote by D_z their employment status when externally assigned $Z = z$. The conditional distributions of $\exp(Y_{z,D_z}) | T = t$ for each $t \in \mathcal{T}$ are assumed to follow Lognormal($\mu_{z|t}, \sigma_{z|t}$) distributions. Here, $\sigma_{z|t} = 1$ for all combinations of z and t , indicating that variability only arises

³³By construction, the outcomes (wages) simulated here are independent of the dataset used in Lee (2009).

through $\mu_{z|t}$ between types. We present two distinct potential earnings distribution models as described in Table 6 and Table 7.

TABLE 6. Design 1

t	Dist. of $\exp(Y_{1,D_1}) T = t$	Dist. of $\exp(Y_{0,D_0}) T = t$
$(0, L)$	Lognormal(9.5, 1)	Lognormal(9.5, 1)
$(0, H)$	Lognormal(11.5, 1)	Lognormal(9.5, 1)
(L, H)	Lognormal(16.5, 1)	Lognormal(9.5, 1)
(H, L)	Lognormal(9.75, 1)	Lognormal(9.6, 1)
(L, L)	Lognormal(9.5, 1)	Lognormal(9.5, 1)
(H, H)	Lognormal(14.5, 1)	Lognormal(14.5, 1)

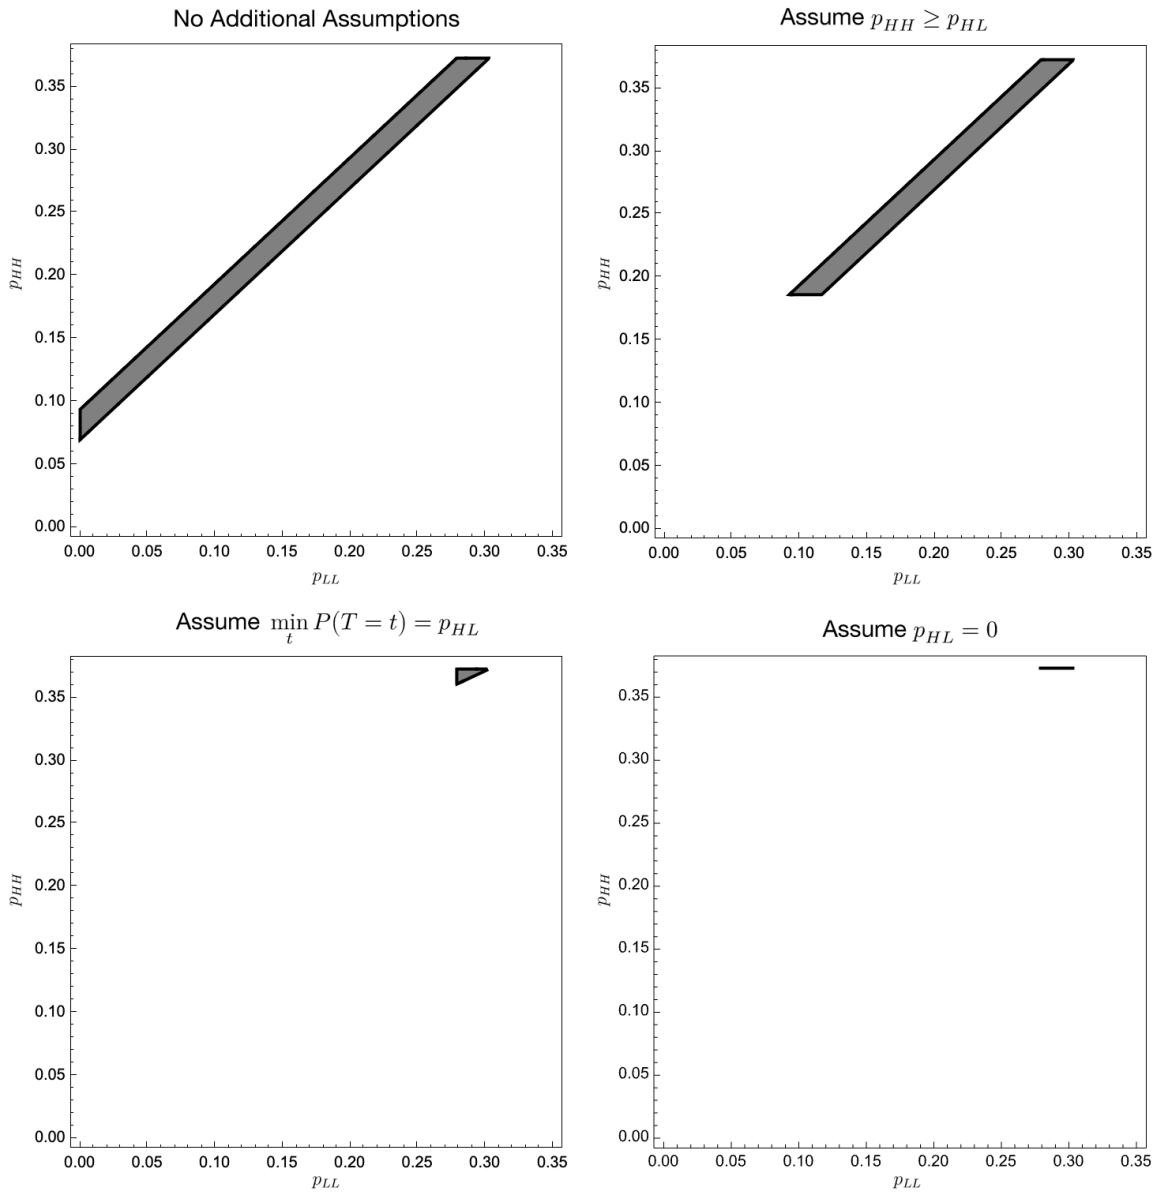
TABLE 7. Design 2

t	Dist. of $\exp(Y_{1,D_1}) T = t$	Dist. of $\exp(Y_{0,D_0}) T = t$
$(0, L)$	Lognormal(10.5, 1)	Lognormal(9.5, 1)
$(0, H)$	Lognormal(12.5, 1)	Lognormal(9.5, 1)
(L, H)	Lognormal(14.5, 1)	Lognormal(9.5, 1)
(H, L)	Lognormal(10.5, 1)	Lognormal(10.5, 1)
(L, L)	Lognormal(10.5, 1)	Lognormal(9.5, 1)
(H, H)	Lognormal(14, 1)	Lognormal(12, 1)

A.0.2. *Simulations Results.* We begin by exploring the geometry of $\Theta_I(\mathcal{R}_T)$, and demonstrate how incorporating further assumptions regarding response types can significantly shrink its shape.

Figures 11 and 13, along with Tables 8 and 9, present the outcomes for designs 1 and 2, respectively. Initially, we compute Lee bounds, which set identifies the total effect $\mathbb{E}(Y_{1,D_1} - Y_{0,D_0}|D_0 > 0, D_1 > 0)$, as in Lemma 2.

Within the framework of design 1, these bounds lie entirely within the positive quadrant, excluding 0. When ignoring firm heterogeneity in wages, these bounds suggest that job training increases wage rates for the always-employed, i.e., $AE = \{(L, L), (H, H), (L, H), (H, L)\}$. However, this DGP is consistent with the fact that the true within-firm causal effect for



Notes: The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed in addition to Assumption 2. In the two bottom panels of the figure, we shrink the scale of the axes to make the identified set visible.

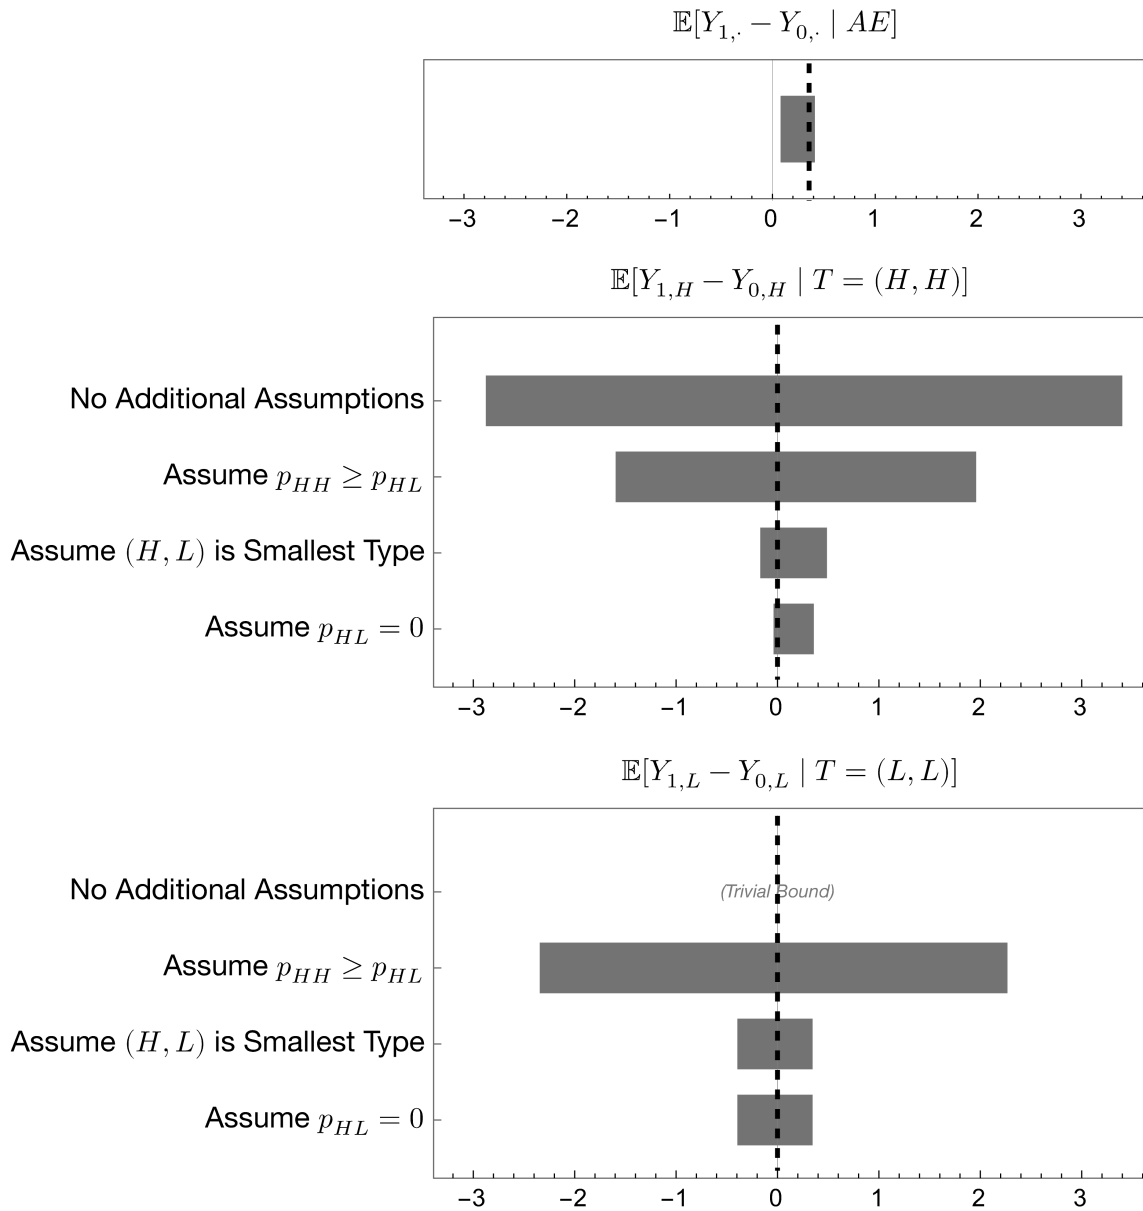
FIGURE 10. Identified set for $(p_{L,L}, p_{H,H})$ in both simulated DGPs.

TABLE 8. Results for Design 1.

Parameter	True value	Assumptions	Bounds	
			Lower	Upper
$E(Y_{1,D_1} - Y_{0,D_0} AE)$	0.3574	Assumptions 1-2	0.0785	0.4131
		Assumptions 1-2	-2.8752	3.4005
		Assumptions 1-2 and $P(T = (H, H)) \geq P(T = (H, L))$	-1.5958	1.9590
$E(Y_{1,H} - Y_{0,H} T = (H, H))$	0.000	Assumptions 1-2 and $\min_{\tau} P(T = \tau) = P(T = (H, L))$	-0.1699	0.4874
		Assumptions 1-2 and $P(T = (H, L)) = 0$	-0.0385	0.3588
		Assumptions 1-2	<i>Trivial Bounds</i>	
$E(Y_{1,L} - Y_{0,L} T = (L, L))$	0.000	Assumptions 1-2 and $\min_{\tau} P(T = \tau) = P(T = (H, L))$	-2.3443	2.2676
		Assumptions 1-2 and $P(T = (H, L)) = 0$	-0.3960	0.3460
		Assumptions 1-2	<i>Trivial Bounds</i>	

$t \in \{(L, L), (H, H), (L, H), (H, L)\}$ is 0, i.e. $\mathbb{E}(Y_{1,H} - Y_{0,H}|T = t) = \mathbb{E}(Y_{1,L} - Y_{0,L}|T = t) = 0 \forall t \in \{(L, L), (H, H), (L, H), (H, L)\}$. That is, the positive effect recovered by Lee bounds purely captures the effect of job training on sorting rather than a change in wages. This distinction is further clarified through the equation below, which illustrates that whenever $\mathbb{E}(Y_{1,H} - Y_{0,H}|T = (H, H)) = \mathbb{E}(Y_{1,L} - Y_{0,L}|T = (L, L)) = 0$, the observed changes in wages captured by Lee's bounds are primarily a consequence of sorting induced by job training:


$$\begin{aligned} \mathbb{E}(Y_{1,D_1} - Y_{0,D_0}|D_0 > 0, D_1 > 0) &= \frac{p_{L,H}}{\mathbb{P}(AE)} \mathbb{E}[Y_{1,H} - Y_{0,L}|T = (L, H)] \\ &\quad + \frac{p_{H,L}}{\mathbb{P}(AE)} \mathbb{E}[Y_{1,L} - Y_{0,H}|T = (H, L)]. \end{aligned}$$



Notes: Dashed black line indicates the true value of the parameter.

FIGURE 11. Lee (2009) Bounds and Multilayered Bounds in Design 1.

Applying our bounds on the within-firm effects, we can provide a more accurate assessment of the impact of job training on wages. Our bounds for $\mathbb{E}(Y_{1,H} - Y_{0,H} | T = (H, H))$ and $\mathbb{E}(Y_{1,L} - Y_{0,L} | T = (L, L))$ always include 0 across the different scenarios which correspond



Notes: Thick dashed black line indicates true value of the aggregate treatment effect. Shaded areas indicate sharp bounds and thin dashed lines indicate the result of the “naïve approach” of taking the weighted average of firm-level bounds (see Remark 2).

FIGURE 12. Aggregate Multilayered Bounds in Design 1.

to different assumptions on response types, suggesting that there is not enough evidence in the data to support the hypothesis that job training has a causal effect on wages.

In design 2, the Lee bounds also lie entirely within the positive quadrant, excluding 0. However, in this case, our DGP is consistent with the fact that the true within-firm causal effects for $t \in \{(L, L), (H, H), (L, H), (H, L)\}$ are strictly positive. Interestingly, our bounds for $\mathbb{E}(Y_{1,H} - Y_{0,H} | T = (H, H))$ and $\mathbb{E}(Y_{1,L} - Y_{0,L} | T = (L, L))$ lie entirely within the positive quadrant, excluding 0, when restricting the response types. This shows that our bounds are informative enough to reveal that job training directly influences wages.

APPENDIX B. ADDITIONAL RESULTS

B.1. Generalization of Horowitz and Manski (1995) Bounds.

TABLE 9. Results of Design 2.

Parameter	True value	Assumptions	Bounds	
			Lower	Upper
$E(Y_{1,D_1} - Y_{0,D_0} AE)$	1.7472	Assumptions 1-2	1.5435	1.8029
		Assumptions 1-2	-0.8958	4.9802
		Assumptions 1-2 and $P(T = (H, H)) \geq P(T = (H, L))$	0.3698	3.7244
$E(Y_{1,H} - Y_{0,H} T = (H, H))$	2.000	Assumptions 1-2 and $\min_{\tau} P(T = \tau) = P(T = (H, L))$	1.7503	2.3440
		Assumptions 1-2 and $P(T = (H, L)) = 0$	1.8741	2.2194
		Assumptions 1-2	<i>Trivial Bounds</i>	
$E(Y_{1,L} - Y_{0,L} T = (L, L))$	1.000	Assumptions 1-2 and $\min_{\tau} P(T = \tau) = P(T = (H, L))$	-1.3443	3.2676
		Assumptions 1-2 and $P(T = (H, L)) = 0$	0.6040	1.3460
		Assumptions 1-2	0.6040	1.3460

Lemma 5. Consider a vector of random variables (W, W_1, \dots, W_K) satisfying the following condition for all w :

$$F_W(w) = \sum_{k=0}^K \gamma_k F_{W_k}(w) \quad (\text{B.1})$$


where $\gamma_k \geq 0$, $F_{W_k}(\cdot)$ represents the cumulative distribution function (CDF) of W_k , and $F_W(\cdot)$ represents the CDF of W .

The results are as follows:

- (i) The bounds for a fixed component of the mixture are point-wise sharp and given by:

$$\mathbb{E}[F_W^{-1}(U)|U \leq \gamma_k] \leq \mathbb{E}[W_k] \leq \mathbb{E}[F_W^{-1}(U)|U \geq 1 - \gamma_k], \quad (\text{B.2})$$

and this for $k \in \{1, \dots, K\}$ where $U \sim \text{Uniform}[0, 1]$.




Notes: Dashed black line indicates the true value of the parameter.

FIGURE 13. Lee (2009) Bounds and Multilayered Bounds in Design 2.

- (ii) *The bounds for a weighted average of fixed components of the mixture are point-wise sharp and given by:*

$$\mathbb{E} \left[F_W^{-1}(U) | U \leq \sum_{k=l}^{l'} \gamma_k \right] \leq \sum_{k=l}^{l'} \frac{\gamma_k}{\sum_{k=l}^{l'} \gamma_k} \mathbb{E}[W_k] \leq \mathbb{E} \left[F_W^{-1}(U) | U \geq 1 - \sum_{k=l}^{l'} \gamma_k \right], \quad (\text{B.3})$$



Notes: Thick dashed black line indicates true value of the aggregate treatment effect. Shaded areas indicate sharp bounds and thin dashed lines indicate the result of the “naïve approach” of taking the weighted average of firm-level bounds (see Remark 2).

FIGURE 14. Aggregate Multilayered Bounds in Design 2.

Lemma 5(i) extends the Horowitz and Manski (1995) bounds to scenarios involving mixtures with more than two components. Lemma 5(ii) introduces novel bounds on the weighted average of certain mixture components spanning l to l' . It is noteworthy that

$$\mathbb{E} \left[F_W^{-1}(U) | U \leq \sum_{k=l}^{l'} \gamma_k \right] \geq \sum_{k=l}^{l'} \frac{\gamma_k}{\sum_{k=l}^{l'} \gamma_k} \mathbb{E}[F_W^{-1}(U) | U \leq \gamma_k],$$

indicating that utilizing the point-wise bounds from Lemma 5(i) (Horowitz and Manski bounds) to establish bounds on the weighted average incurs some information loss.

B.2. Bounds on Aggregate Local Control Direct Effect (LCDE). To derive bounds on the aggregate LCDE, defined as $\sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d'=l}^{l'} p_{d',d'}} \text{LCDE}(d | d, d)$, one might be tempted to adopt a naive approach by taking a weighted average of the pointwise sharp bounds derived in 1 (i). However, this approach not only fails to provide sharp bounds on the aggregate quantity but may also yield invalid bounds.

The primary issue lies in the fact that both the weights and the LCDE quantities are only set-identified. Consequently, the naive aggregation does not preserve the sharpness or validity of the bounds. To better illustrate this problem, let us provide more details:

- (1) The sharp bounds on the aggregate LCDE can be tighter than the weighted average of the marginal bounds on $\text{LCDE}(d \mid d, d)$ for $d = l, \dots, l'$ when these marginal bounds are achieved at values of $p_{d,d}$ for $d = l, \dots, l'$ that are not *jointly* attainable. Specifically, this occurs whenever $(\underline{p}_{d,d} : d = l, \dots, l')$ does not lie within the identified set for $(p_{d,d} : d = l, \dots, l')$. As a result, the sharp bounds, which are obtained by optimizing over the identified set for the *vector* $(p_{d,d} : d = l, \dots, l')$, may be tighter.
- (2) The sharp bounds on the LCDE can be wider than the weighted average of the marginal LCDE bounds when there exist $d, d' \in \{l, \dots, l'\}$ such that the lower bound for $\text{LCDE}(d \mid d, d)$ is significantly smaller than the lower bound for $\text{LCDE}(d' \mid d', d')$, and $\underline{p}_{d,d}$ is also much smaller than $\underline{p}_{d',d'}$.

In such cases, the choice of $p_{d,d}$ directly affects the relative weight of $\text{LCDE}(d \mid d, d)$ in the aggregate LCDE. Consequently, the value of the objective function in the optimization problem defining the lower bound in Theorem 1 (iii) may actually decrease as $p_{d,d}$ increases.

To illustrate this more clearly, consider a simple example with only two groups, a and b . For each $d \in \{a, b\}$, let $L_d(p)$ denote the sharp lower bound for $\text{LCDE}(d \mid d, d)$ under the assumption that $p_{d,d} = p$. Suppose $\underline{p}_{b,b} > \underline{p}_{a,a} = 0$, so that the lower bound for $\text{LCDE}(a \mid a, a)$, given by $L_a(\underline{p}_{a,a}) = L_a(0) = y_L$, is the trivial bound. Additionally, assume the lower bound for $\text{LCDE}(b \mid b, b)$ is non-trivial, i.e., $L_b(\underline{p}_{b,b}) > y_L$.

In this scenario, the weighted average of these lower bounds simplifies to:

$$\frac{\underline{p}_{a,a} L_a(\underline{p}_{a,a}) + \underline{p}_{b,b} L_b(\underline{p}_{b,b})}{\underline{p}_{a,a} + \underline{p}_{b,b}} = L_b(\underline{p}_{b,b}),$$

since $\underline{p}_{a,a} = 0$.

Now, consider the case where there exists a point $p > 0$ such that $(p, \underline{p}_{b,b})$ lies within the identified set for $(p_{a,a}, p_{b,b})$, and $L_a(p) < L_b(\underline{p}_{b,b})$. In this case, we have:

$$\frac{p L_a(p) + \underline{p}_{b,b} L_b(\underline{p}_{b,b})}{p + \underline{p}_{b,b}} < L_b(\underline{p}_{b,b}).$$

By definition, the sharp lower bound given in Theorem 1 (iii) will be at least as small as the left-hand side above, resulting in a smaller lower bound for the aggregate

LCDE. Similarly, it is also possible for the sharp upper bound to be larger than the weighted average of the marginal upper bounds.

APPENDIX C. PROOFS OF THE MAIN RESULTS

C.1. Proof of Lemma 2.

Proof. Note that (ii) and (iii) follow immediately from (i), so it suffices to show (i). First, notice that

$$\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] = \mathbb{E}[Y_{1,D_1} - Y_{0,D_1} | D_0 > 0, D_1 > 0] \quad (\text{C.1})$$

$$+ \mathbb{E}[Y_{0,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] . \quad (\text{C.2})$$

Next, we have that

$$\begin{aligned} & \mathbb{E}[Y_{1,D_1} - Y_{0,D_1} | D_0 > 0, D_1 > 0] \\ &= \sum_{d=1, d'=1}^{K, K} \mathbb{E}[Y_{1,D_1} - Y_{0,D_1} | D_0 = d', D_1 = d] \mathbb{P}(D_0 = d', D_1 = d | D_0 > 0, D_1 > 0) \\ &= \sum_{d=1, d'=1}^{K, K} \mathbb{E}[Y_{1,d} - Y_{0,d} | D_0 = d', D_1 = d] \mathbb{P}(D_0 = d', D_1 = d | D_0 > 0, D_1 > 0) \\ &= \sum_{d=1, d'=1}^{K, K} LCDE(d|d', d) \mathbb{P}(T = (d', d) | D_0 > 0, D_1 > 0) , \end{aligned}$$

and similarly,

$$\begin{aligned} & \mathbb{E}[Y_{0,D_1} - Y_{0,D_0} | D_0 > 0, D_1 > 0] \\ &= \sum_{d=1, d'=1}^{K, K} \mathbb{E}[Y_{0,D_1} - Y_{0,D_0} | D_0 = d', D_1 = d] \mathbb{P}(D_0 = d', D_1 = d | D_0 > 0, D_1 > 0) \\ &= \sum_{d=1, d'=1}^{K, K} \mathbb{E}[Y_{0,d} - Y_{0,d'} | D_0 = d', D_1 = d] \mathbb{P}(D_0 = d', D_1 = d | D_0 > 0, D_1 > 0) \\ &= \sum_{d=1, d'=1, d \neq d'}^{K, K} \mathbb{E}[Y_{0,d} - Y_{0,d'} | D_0 = d', D_1 = d] \mathbb{P}(D_0 = d', D_1 = d | D_0 > 0, D_1 > 0) \\ &= \sum_{d=1, d'=1, d \neq d'}^{K, K} LCIE(0, d, d' | d', d) \mathbb{P}(T = (d', d) | D_0 > 0, D_1 > 0) , \end{aligned}$$

and plugging these values back into (C.1) immediately implies the result. \square

C.2. Proof of Lemma 4.

Proof. By the definition of T , it is clear that the response-type probabilities satisfy the eqs (4.4, 4.5). It suffices to show that given any solution $(\mathbf{p}_{(d,d')} : (d', d) \in \{0, \dots, K\}^2) \geq 0$ such that $\sum_{d=0, d'=0}^{K, K} \mathbf{p}_{(d,d')} = 1$ and satisfying, for every $d \in \{0, \dots, K\}$,

$$\mathbb{P}(D = d | Z = 1) = \sum_{d'=0}^K \mathbf{p}_{(d,d')} \quad (\text{C.3})$$

$$\text{and } \mathbb{P}(D = d | Z = 0) = \sum_{d'=0}^K \mathbf{p}_{(d,d')} , \quad (\text{C.4})$$

there exists a joint distribution Q of $((Y_{0,d} : d \in \{0, \dots, K\}), (Y_{1,d} : d \in \{0, \dots, K\}), D_0, D_1, Z)$ such that

$$\left(\mathbb{P}_Q [T = (d, d')] : (d', d) \in \{0, \dots, K\}^2 \right) = \left(\mathbf{p}_{(d,d')} : (d', d) \in \{0, \dots, K\}^2 \right) ,$$

and Q induces a distribution of (Y, D, Z) under (2.5, 2.6) and Assumption 1, that is consistent with the observed data. Since Y is not observed when $D = 0$, set $Y | D = 0, Z = 1$ and $Y | D = 0, Z = 0$ to arbitrary distributions, so that we can treat (Y, D, Z) as observed. We will now construct a Q that induces this distribution.

Define $\mathcal{Y}_{z,d} := \text{supp}(Y | D = d, Z = z)$ and, for each $d \in \{0, \dots, K\}$, $z \in \{0, 1\}$, and $(d', d'') \in \{0, \dots, K\}^2$ define the CDF $F_{(d',d'')}^{(z,d)}$ as

$$F_{(d',d'')}^{(z,d)}(y) = P(Y \leq y | D = d, Z = z) ,$$

and note that it does not depend on (d', d'') . Next, for every $(d', d'') \in \{0, \dots, K\}^2$, let $C_{(d',d'')}$ be an arbitrary copula of dimension $|\{0, 1\} \times \{0, \dots, K\}|$. Finally, define Q as

$$Q(y, t, z) := C_t \left(\left(F_t^{(z,d)}(y_{(z,d)}) : (z, d) \in \{0, 1\} \times \{0, \dots, K\} \right) \right) \times \mathbf{p}_{(t)} \times P(Z = z) ,$$

for any $y \in \prod_{(z,d) \in \{0,1\} \times \{0,\dots,K\}} \mathcal{Y}_{z,d}$, $t \in \{0, \dots, K\}^2$, and $z \in \{0, 1\}$, where $Q(y, t, z)$ is shorthand for

$$Q((Y_{z,d} : (z, d) \in \{0, 1\} \times \{0, \dots, K\}) \leq y, (D_0, D_1) = t, Z = z)$$

Next, for all $(d', d'') \in \{0, \dots, K\}^2$, let the conditional joint distribution

$$Q(Y_{z,d} \leq y_{z,d} | T = (d', d'')) .$$

By construction, Q satisfies Assumption 1 and $\left(\mathbb{P}_Q [T = (d, d')] : (d', d) \in \{0, \dots, K\}^2\right) = \left(\mathbf{p}_{(d, d')} : (d', d) \in \{0, \dots, K\}^2\right)$. The result now follows immediately by noting that Q induces the observed data distribution under (2.5, 2.6) since, for any $y \in \mathcal{Y}_{z, d}$, $d \in \{0, \dots, K\}$ and $z \in \{0, 1\}$, we have that

$$\begin{aligned}
Q(Y_{z, d} \leq y, D_z = d, Z = z) &= Q(Y_{z, d} \leq y | D_z = d, Z = z) Q(D_z = d | Z = z) Q(Z = z) \\
&= Q(Y_{z, d} \leq y | D_z = d) Q(D_z = d) P(Z = z) \\
&= P(Z = z) \sum_{d'=1}^{K, K} Q(Y_{z, d} \leq y | D_z = d, D_{1-z} = d') \left((1-z)\mathbf{p}_{(d, d')} + z\mathbf{p}_{(d', d)} \right) \\
&= P(Z = z) \sum_{d'=1}^{K, K} \left((1-z)F_{(d, d')}^{(z, d)}(y) \mathbf{p}_{(d, d')} + zF_{(d', d)}^{(z, d)}(y) \mathbf{p}_{(d', d)} \right) \\
&= P(Z = z) P(Y \leq y | D = d, Z = z) \sum_{d'=1}^{K, K} \left((1-z)\mathbf{p}_{(d, d')} + z\mathbf{p}_{(d', d)} \right) \\
&= P(Z = z) P(Y \leq y | D = d, Z = z) P(D = d | Z = z) \\
&= P(Y \leq y, D = d, Z = z)
\end{aligned}$$

where the second equality follows from Q satisfying Assumption 1, and the penultimate equality follows from \mathbf{p} satisfying (C.3) and (C.4). \square

C.3. Proof of Lemma 4.

Proof. Given Lemma 3 above, this result follows immediately from Theorem 3 in Vayalinkal (2024). Since our proof of Lemma 3 is constructive, however, we can also argue directly, as follows. Note that both parts below proceed by showing the contrapositive.

(\Leftarrow) If Assumption 1 and \mathcal{R}_T are consistent with the data, then there exists a joint distribution Q of $((Y_{0, d} : d \in \{0, \dots, K\}), (Y_{1, d} : d \in \{0, \dots, K\}), D_0, D_1, Z)$ that is consistent with the observed data distribution such that the response-type probabilities induced by Q is in $\Theta_I(\mathcal{R}_T)$ and so $\Theta_I(\mathcal{R}_T) \neq \emptyset$.

(\Rightarrow) Suppose that $\Theta_I(\mathcal{R}_T) \neq \emptyset$, then there exists $\mathbf{p} = \left(\mathbf{p}_{(d, d')} : (d', d) \in \{0, \dots, K\}^2\right) \in \Theta_I(\mathcal{R}_T)$. Now, our proof of Lemma 3 shows that we can construct a joint distribution Q of $((Y_{0, d} : d \in \{0, \dots, K\}), (Y_{1, d} : d \in \{0, \dots, K\}), D_0, D_1, Z)$ such that Q satisfies Assumption 1 and induces the observed data distribution under (2.5, 2.6), and

$$\left(\mathbb{P}_Q [T = (d, d')] : (d', d) \in \{0, \dots, K\}^2\right) = \left(\mathbf{p}_{(d, d')} : (d', d) \in \{0, \dots, K\}^2\right),$$

which implies that Q also satisfies \mathcal{R}_T , as required, since $\mathbf{p} \in \Theta_I(\mathcal{R}_T)$. \square

C.4. Proof of Lemma 5.

Proof. Since (i) is just a special case of (ii), it suffices to show (ii). Let W and $\{W_k\}_{k=0}^K$ be defined on the common probability space $(\Omega, \Sigma, \mathbb{P})$, and take values in the set $\mathcal{W} \subseteq \mathbb{R}$, equipped with the Borel sigma algebra and a probability measure μ . Moreover, let μ be such that W and $\{W_k\}_{k=0}^K$ are μ -integrable and have densities with respect to μ . Denote the μ -density of W by f_W , and denote the μ -density of W_k by f_{W_k} , for $k \in \{0, \dots, K\}$.

First, we show that the bounds are valid. Define $\bar{\gamma} := \sum_{k=l}^{l'} \gamma_k$ and let $U \sim \text{Uniform}[0, 1]$. Now, suppose that

$$\mathbb{E} [F_W^{-1}(U) | U \leq \bar{\gamma}] > \sum_{k=l}^{l'} \frac{\gamma_k}{\bar{\gamma}} \mathbb{E}[W_k],$$

then there must exist w such that

$$\mathbb{P} (F_W^{-1}(U) \leq w | U \leq \bar{\gamma}) < \sum_{k=l}^{l'} \frac{\gamma_k}{\bar{\gamma}} \int_{(-\infty, w]} f_{W_k}(x) d\mu(x) = \frac{1}{\bar{\gamma}} \int_{(-\infty, w]} \sum_{k=l}^{l'} \gamma_k f_{W_k}(x) d\mu(x) .$$

Now, note that

$$\begin{aligned} \mathbb{P} (F_W^{-1}(U) \leq w | U \leq \bar{\gamma}) &= \mathbb{P} (F_W^{-1}(\bar{\gamma}U) \leq w) = \mathbb{P} (\bar{\gamma}U \leq F_W(w)) \\ &= \frac{F_W(w)}{\bar{\gamma}} = \frac{\int_{(-\infty, w]} f_W(x) d\mu(x)}{\bar{\gamma}}, \end{aligned}$$

but this implies that

$$\int_{(-\infty, w]} \sum_{k=0}^K \gamma_k f_{W_k}(x) d\mu(x) = \int_{(-\infty, w]} f_W(x) d\mu(x) < \int_{(-\infty, w]} \sum_{k=l}^{l'} \gamma_k f_{W_k}(x) d\mu(x) ,$$

which, in turn, implies that

$$\int_{(-\infty, w]} \sum_{k=0}^{l-1} \gamma_k f_{W_k}(x) + \sum_{k=l'+1}^K \gamma_k f_{W_k}(x) d\mu(x) < 0 ,$$

a contradiction. Therefore, the lower bound is valid. The validity of the upper bound follows from the analogous argument.

Define $\underline{\gamma} := \bar{\gamma} + \sum_{k'=0}^{l-1} \gamma_{k'}$. Now, sharpness of the lower bound follows immediately by defining each F_{W_k} as follows

$$F_{W_k}(w) := \begin{cases} \mathbb{P}\left(F_{W_k}^{-1}(U) \leq w | U \in \left(\sum_{k'=l}^{k-1} \gamma_{k'}, \sum_{k'=l}^k \gamma_{k'}\right)\right) & \text{if } k \in \{l, \dots, l'\} \\ \mathbb{P}\left(F_{W_k}^{-1}(U) \leq w | U \in \left(\bar{\gamma} + \sum_{k'=0}^{k-1} \gamma_{k'}, \bar{\gamma} + \sum_{k=0}^k \gamma_{k'}\right)\right) & \text{if } k \in \{0, \dots, l-1\} \\ \mathbb{P}\left(F_{W_k}^{-1}(U) \leq w | U \in \left(\underline{\gamma} + \sum_{k'=l'+1}^{k-1} \gamma_{k'}, \underline{\gamma} + \sum_{k'=l'+1}^k \gamma_{k'}\right)\right) & \text{all other } k \end{cases} .$$

Sharpness of the upper bound follows from the analogous construction. \square

C.5. Proof of Theorem 1.

Proof. Define

$$\mathbf{f}_{(z,d)|d,d'}(y) \equiv \mathbb{P}[T = (d', d)] f_{Y_{z,d}|T}(y|d', d) ,$$

and also define the shorthand notation

$$\mathbf{f}_{z|d,d'}(y) \equiv \mathbb{P}[T = (d', d)] f_{Y_{z,T(z)}|T}(y|d', d) .$$

Consider the vector of weighted response-type conditional densities

$$\mathbf{f}(y) \equiv (\mathbf{f}_{(z,d'')|d,d'}(y) : d, d', d'' \in \{0, \dots, K\}, z \in \{0, 1\}) .$$

First, note that \mathbf{f} must satisfy (4.2) and (4.3), i.e. we must have that

$$\begin{aligned} f_{Y,D=d|Z=1}(y) &= \sum_{d'=0}^K \mathbb{P}[T = (d', d)] f_{Y_{1,d}|T}(y|d', d) = \sum_{d'=0}^K \mathbf{f}_{1|d',d}(y) \\ f_{Y,D=d|Z=0}(y) &= \sum_{d'=0}^K \mathbb{P}[T = (d, d')] f_{Y_{0,d}|T}(y|d, d') = \sum_{d'=0}^K \mathbf{f}_{0|d,d'}(y) \end{aligned}$$

for all $d \in \{1, \dots, K\}$. Second, we must also have that there exists a $p := (p_t : t \in \text{supp}(T)) \in \Theta_I(\mathcal{R}_T)$ such that $\int_{y_L}^{y_U} \mathbf{f}_{(z,d'')|d',d}(y) d\mu(y) = p_{d',d}$ for all $z \in \{0, 1\}$ and $d, d', d'' \in \{0, \dots, K\}$. Finally, we must have that each component of \mathbf{f} is a non-negative function supported on (a subset of) $[y_L, y_U]$

The remainder of the proof proceeds in two parts. We first show that these conditions are sharp (i.e. they define the identified set of \mathbf{f}). We then show that this allows us to complete the proof using the sharp bounds on mixture components given in Lemma 5.

A Preliminary: Identified set of \mathbf{f} . Note that for any type t and $z \in \{0, 1\}$, $f_{Y_{z,d''}|T}(y|t)$ is independent of the data whenever $d'' \neq t(z)$, and so, is only constrained to be a density that is supported over (a subset of) $[y_L, y_U]$, the support of $Y_{z,d''}$; this immediately implies that the sharp identification region for the expectation of any such component is simply

$[y_L, y_U]$. Given Lemma 3 above, it now follows from Theorem 3.2 in Vayalinkal (2024) that these conditions are sharp, i.e. any \mathbf{f} satisfying these conditions is consistent with the data. We summarize the argument here, as follows.

First, note that the observed data depends only on the (i) the distribution of Z (F_Z), (ii) the marginal distribution of D_z for each $z \in \{0, 1\}$, and (iii) the conditional marginal distribution of $Y_{z,d}$ given $D_z = d, Z = z$ for all $d \in \{1, \dots, K\}$ and $z \in \{0, 1\}$. For any joint distribution $((Y_{z,d} : d \in \{0, \dots, K\}, z \in \{0, 1\}), T, Z) \sim Q$, let \mathbf{f}_Q be the vector of weighted response-type conditional densities implied by Q . Given \mathbf{f} satisfying the conditions above, we construct a Q with $\mathbf{f}_Q = \mathbf{f}$ as follows: define $Q_Z = F_Z$, $Q(T = t) = \int_{\mathbf{y}} \mathbf{f}_{1|t}(y) d\mu(y)$, and define

$$\begin{aligned} & Q(Y_{0,0} \leq y_{0,0}, \dots, Y_{0,K} \leq y_{0,K}, Y_{1,0} \leq y_{1,0}, \dots, Y_{1,K} \leq y_{1,K}, T = t, Z = z) \\ &= \left(\prod_{k=0}^K \int_{y_L}^{y_{0,k}} \mathbf{f}_{(0,k)|t}(y) d\mu(y) \right) \left(\prod_{k=0}^K \int_{y_L}^{y_{1,k}} \mathbf{f}_{(0,k)|t}(y) d\mu(y) \right) Q(T = t) Q(Z = z) . \end{aligned}$$

The above construction assumes that the potential outcome distributions are independent given T , but any dependence structure (copula) can be used, after conditioning on a value of T . Suppose we are given a Q such that \mathbf{f}_Q satisfies the conditions above. By construction, $Q_Z = F_Z$, $Q(D_z = d) = \sum_{t:t(z)=d} Q(T = t) = \sum_{t:t(z)=d} \int_{y_L}^{y_U} \mathbf{f}_{z|t}(y) d\mu(y) = P(D = d|Z = z)$ for all $d \in \{1, \dots, K\}$ and $z \in \{0, 1\}$. This also implies $Q(D_z = 0) = P(D = 0|Z = z)$ by the definition of Θ_I and, finally, we have that for any $z \in \{0, 1\}, d \in \{1, \dots, K\}$

$$\begin{aligned} Q(Y_{z,d} \leq y | D_z = d, Z = z) &= \sum_{t:t(z)=d} \int_{y_L}^y \mathbf{f}_{z|t}(y) d\mu(y) = \int_{y_L}^y \sum_{t:t(z)=d} \mathbf{f}_{z|t}(y) d\mu(y) \\ &= \int_{y_L}^y f_{Y,D=d|Z=z}(y) d\mu(y) = P(Y \leq y | D = d, Z = z) , \end{aligned}$$

as required, showing that the above conditions define the identified set for \mathbf{f} .

Remainder of proof. Since the two mixtures given by (4.2) and (4.3) do not share any components, the above result reduces the problem of finding bounds on the conditional expectation of $Y_{d,z}$ given T to the problem of finding sharp bounds on expectations of mixture components. Therefore, we now complete the proof using the results given in Lemma 5, as follows.

Proof of (i) and (ii): For any type (d', d) and any $z \in \{0, 1\}$, the weighted conditional density $\mathbb{P}[T = (d', d)] \times f_{Y_{z,T(z)}|T}(y|d', d)$ only appears in at most one of (4.2) and (4.3). Therefore, sharp bounds on the expectation of any such component can be obtained as the bounds of Horowitz and Manski (1995) (HM), which are the bounds provided in Lemma 5(i)

evaluated at the smallest feasible value of γ_k . This immediately implies the validity and sharpness of (ii) and the last two parts of (i) above. This also implies that the bounds in the first part of (i) are valid, as follows: (I) $\text{LCDE}(d|d, d)$ consists of the difference in expectation of two such components, and (II) the lower (upper) bound is given by the HM lower (upper) bound of the first component minus the HM upper (lower) bound of the second component. For sharpness, first note that $\text{LCDE}(d|d, d)$ consists of the difference in expectation of two components, each of which belongs to a *different* mixture of the two defined by (4.2) and (4.3). Since these two mixtures do not share any components (only weights), the two HM bounds can be attained jointly whenever the weights $\underline{\gamma}_{d,d}^{1,r}$ and $\underline{\gamma}_{d,d}^{0,r}$ are jointly feasible. The result now follows by noting that $\underline{\gamma}_{d,d}^{1,r}$ and $\underline{\gamma}_{d,d}^{0,r}$ are jointly feasible if and only if $\mathbb{P}(D = d|Z = 1)\underline{\gamma}_{d,d}^{1,r} = \mathbb{P}(D = d|Z = 0)\underline{\gamma}_{d,d}^{0,r} = \underline{p}_{d,d'}^r$ belongs to the identified set for $p_{d,d'}$ which is true by definition of $\underline{p}_{d,d'}^r$.

Proof of (iii): Finally, for (iii), note that

$$\begin{aligned} & \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \text{LCDE}(d|d, d) = \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \mathbb{E}[Y_{1,d} - Y_{0,d}|T = (d, d)] \\ &= \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \mathbb{E}[Y_{1,d}|T = (d, d)] - \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \mathbb{E}[Y_{0,d}|T = (d, d)] \\ &= \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \int_{\mathcal{Y}} y \mathbf{f}_{1|d,d}(y) d\mu(y) - \sum_{d=l}^{l'} \frac{p_{d,d}}{\sum_{d=l}^{l'} p_{d,d}} \int_{\mathcal{Y}} y \mathbf{f}_{0|d,d}(y) d\mu(y) , \end{aligned}$$

where the first term is the expectation of the aggregation of components of the mixture (4.2) and the second term is the expectation of the aggregation of components of the mixture (4.3); we can now use the same argument as above, but now based on Lemma 5(ii), as follows. First, suppose that $\{(p_{d,d} : l \leq d \leq l') \mid p \in \Theta_I(\mathcal{R}_T)\}$ is a singleton, so that the weights are known and there is no optimization required: sharpness now follows immediately from Lemma 5(ii) since the two mixtures do not share any components. Finally, when this is not the case, the sharp bounds are obtained by maximizing (resp. minimizing) the pointwise (in $(p_{d,d} : l \leq d \leq l')$) upper (resp. lower) bound over the identified set for $(p_{d,d} : l \leq d \leq l')$, which is exactly the bounds given in (iii). \square

APPENDIX D. ADDITIONAL EMPIRICAL TABLES AND FIGURES FOR JOB CORPS RCT

D.1. Summary Statistics. Table 10 presents summary statistics for the sample of individuals who have non-missing values for weekly earnings and hours for every week following random assignment. Means and standard deviations for a number of baseline and post-randomization variables are reported separately by treatment status. Consistent with successful randomization in the National Job Corps Study, the table shows that there are no statistical differences in the means of demographic, education, background and baseline employment/income variables across treatment and control groups. This finding aligns with the previous evaluations of the National Job Corps Study.

Table 10 also shows economically and statistically significant differences in employment and earnings outcomes by treatment status, post randomization. We see that 52 weeks after randomization, treatment group hours and earnings are lower than the control group, but 104 weeks after randomization, treatment group hours and earnings exceed the control group. After 208 weeks, the hours and earnings of the treatment group are approximately 8% and 14% higher, respectively, than those of the control group. These differences are statistically significant and are consistent with previous evaluations of the National Job Corps Study.

D.2. Joint Distribution of Wages and Amenities. Figure 15 presents the mean log wage at week 90 according to whether firms provide amenities. At week 90 (208), mean wages at firms that offer amenities are approximately 15% (18%) higher compared to firms that do not. Figure 16 presents the empirical cumulative distribution functions of log wage by firm type, classifying firms according to the provision of health insurance, for treatment and control groups at week 90. For both treated and control units, the distribution of log wages for firms that provide health insurance stochastically dominates the distribution of log wages for firms that do not. This evidence suggests that firms providing amenities pay higher wages than firms that do not.³⁴ A natural follow-up question is whether Job Corps affects the sorting of workers into amenity-providing firms.


D.3. Differential Sorting of Treatment and Control Workers. Lee (2009) focuses on the potential for job training to affect labor supply along the extensive margin but ignores

³⁴Of course, it is possible that firms pay compensating differentials which causally reduce wages. The evidence presented here shows that the variation between firms dominates the variation within firms. This is consistent with evidence in Lamadon et al. (2022) who show that the high-amenities firms are also the more productive firms.

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Female	0.46	0.50	0.45	0.50	-0.01	0.01
Age at baseline	18.35	2.10	18.44	2.16	0.09	0.05
White, non-Hispanic	0.26	0.44	0.27	0.44	0.00	0.01
Black, non-Hispanic	0.49	0.50	0.49	0.50	0.00	0.01
Hispanic	0.17	0.38	0.17	0.37	-0.00	0.01
Other race/ethnicity	0.07	0.26	0.07	0.26	-0.00	0.01
Never married	0.92	0.28	0.92	0.28	0.00	0.01
Married	0.02	0.15	0.02	0.14	-0.00	0.00
Living together	0.04	0.20	0.04	0.19	-0.00	0.00
Separated	0.02	0.14	0.02	0.15	0.00	0.00
Has child	0.19	0.39	0.19	0.39	-0.00	0.01
Number of children	0.27	0.64	0.27	0.65	0.00	0.01
Education	10.11	1.54	10.11	1.56	0.01	0.03
Mother's education	11.46	2.59	11.48	2.56	0.02	0.06
Father's education	11.54	2.79	11.39	2.85	-0.15	0.08
Ever arrested	0.25	0.43	0.25	0.43	-0.00	0.01
Household income						
<3,000	0.25	0.43	0.25	0.44	0.00	0.01
3,000-6,000	0.21	0.41	0.21	0.40	-0.00	0.01
6,000-9,000	0.11	0.32	0.12	0.32	0.00	0.01
9,000-18,000	0.24	0.43	0.24	0.43	-0.00	0.01
>18,000	0.18	0.39	0.18	0.38	-0.00	0.01
Personal income						
<3,000	0.79	0.41	0.79	0.41	-0.00	0.01
3,000-6,000	0.13	0.34	0.13	0.33	-0.00	0.01
6,000-9,000	0.05	0.21	0.05	0.22	0.01	0.00
>9,000	0.03	0.18	0.03	0.17	-0.00	0.00
At baseline						
Have job	0.19	0.39	0.20	0.40	0.01	0.01
Mths. empl. prev. yr.	3.53	4.24	3.60	4.25	0.07	0.09
Had job, prev. yr.	0.63	0.48	0.63	0.48	0.01	0.01
Earnings, prev. yr.	2810.48	4435.62	2906.45	6401.33	95.97	117.10
Usual hours/week	20.91	20.70	21.82	21.05	0.91	0.45
Usual weekly earn.	102.89	116.46	110.99	350.61	8.10	5.09
Post randomization						
Week 52 hours	17.78	23.39	15.30	22.68	-2.49	0.49
Week 104 hours	21.98	26.08	22.64	26.25	0.67	0.56
Week 156 hours	23.88	26.15	25.88	26.57	2.00	0.56
Week 208 hours	25.83	26.25	27.79	25.74	1.95	0.56
Week 52 earn.	103.80	159.89	91.55	149.28	-12.25	3.33
Week 104 earn.	150.41	210.24	157.42	200.27	7.02	4.42
Week 156 earn.	180.88	224.43	203.71	239.80	22.84	4.94
Week 208 earn.	200.50	230.66	227.91	250.22	27.41	5.11
Total 4 yr. earn.	30006.69	26893.60	30800.41	26437.39	793.72	571.83
Sample size	3599		5546		9145	

Notes: Weekly earnings calculated as the sum of total earnings in a given week and are not conditional on employment (i.e., includes 0s for the unemployed).

TABLE 10. Summary statistics by treatment status




Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 15. Mean $\ln(\text{hourly Wage})$ by firm amenity provision at week 90

an additional margin of labor supply: firm choice. If there is scope for job training to affect worker sorting to firms, sample selection is multilayered. Table 11 presents the probability of working in a firm (conditional on employment) that provides observable amenities, at week 90, according to the status of treatment. The evidence shows that treated individuals are more likely to work at firms with job amenities in all but one case (and we also find that this trend persists across all weeks). This is consistent with the evidence presented in Schochet et al. (2008).³⁵

Taken together, Figures 15 and 16 and Table 11 show: (i) firms that provide amenities pay higher wages than firms that do not and (ii) Job Corps affects the sorting of workers into

³⁵In Appendix D.5 we focus on the amenities we use to classify firm types – health benefits H , retirement/pension benefits R and paid vacation V – and categorize jobs into eight mutually exclusive categories based on the amenities available. Appendix Table 15 presents the distribution of workers by the amenity category in which their job falls at week 90. Treated workers are approximately 15% more likely to work in jobs that offer all amenities and 10% less likely to work in jobs that offer no amenities. Reinforcing the finding from Figure 15, hourly wages are approximately 22% higher in jobs with all amenities compared to jobs without amenities. These trends persist qualitatively throughout all weeks.



(a) control units

(b) treated units

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 16. Cumulative distribution function by firm type at week 90, amenity=health

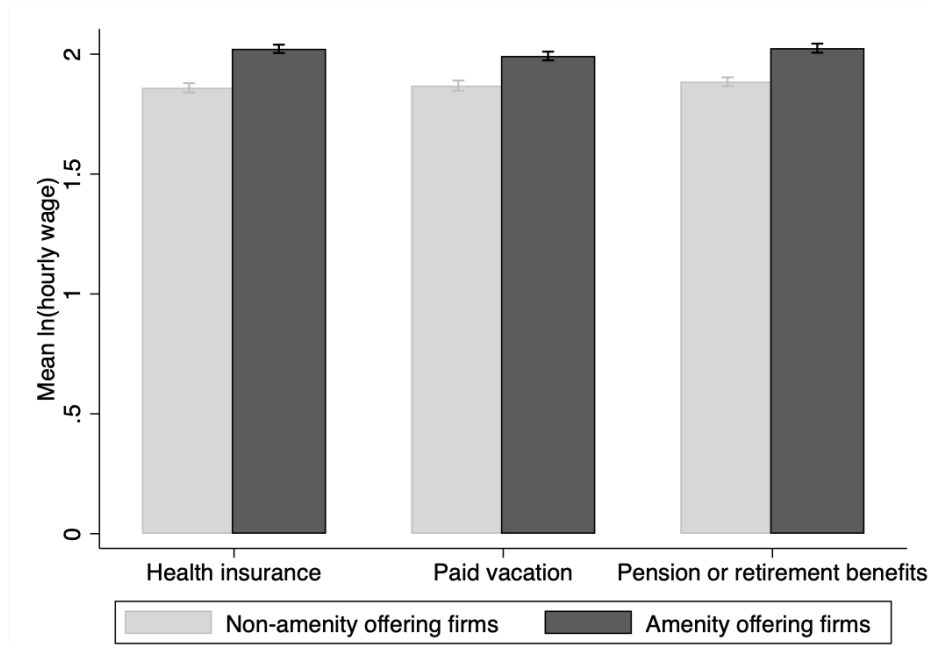
	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Health insurance	0.4860	0.5000	0.5072	0.5001	0.0211	0.0186
Paid sick leave	0.3922	0.4884	0.4251	0.4945	0.0329	0.0183
Paid vacation	0.5407	0.4985	0.5800	0.4937	0.0393	0.0180
Childcare assistance	0.1311	0.3376	0.1422	0.3494	0.0111	0.0127
Flexible hours	0.5330	0.4991	0.5568	0.4969	0.0237	0.0183
Employer-provided transportation	0.1973	0.3981	0.1876	0.3905	-0.0097	0.0147
Pension or retirement benefits	0.3670	0.4822	0.3938	0.4887	0.0268	0.0180
Dental plan	0.3863	0.4871	0.4288	0.4950	0.0425	0.0182
Tuition reimbursement	0.2212	0.4152	0.2602	0.4389	0.0390	0.0158
Employment	0.4368	0.4961	0.4391	0.4963	0.0022	0.0111
Sample size	3288		5091		8379	

Notes: Control and treatment probabilities have interpretation as $P[D = H | D > 0, Z = z]$ for $z \in \{0, 1\}$, respectively, when classifying firms as type H if they provide a given amenity.

TABLE 11. Probability of working at amenity-providing firm at week 90 (conditional on employment)


amenity-providing firms. We therefore conclude that Job Corps training not only affects labor supply along the extensive margin but also along the additional margin of firm choice. As a result, sample selection is multilayered motivating the use of our bounds.

D.4. **Additional descriptive figures and tables.** Appendix D.4 provides descriptive figures and tables from Section 5 for weeks 135, 180 and 208.




Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 17. Mean ln(hourly Wage) by firm amenity provision at week 135




Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 18. Mean $\ln(\text{hourly Wage})$ by firm amenity provision at week 180



Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 19. Mean $\ln(\text{hourly Wage})$ by firm amenity provision at week 208




(a) control units

(b) treated units

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 20. Cumulative distribution function by firm type at week 135, amenity=health




(a) control units

(b) treated units

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 21. Cumulative distribution function by firm type at week 180, amenity=health



(a) control units

(b) treated units

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 22. Cumulative distribution function by firm type at week 208, amenity=health

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Health insurance	0.5375	0.4988	0.5514	0.4975	0.0139	0.0177
Paid sick leave	0.4614	0.4987	0.4552	0.4981	-0.0061	0.0177
Paid vacation	0.6067	0.4887	0.5954	0.4909	-0.0114	0.0171
Childcare assistance	0.1254	0.3313	0.1411	0.3482	0.0157	0.0120
Flexible hours	0.5600	0.4966	0.5794	0.4938	0.0194	0.0174
Employer-provided transportation	0.1899	0.3924	0.1867	0.3898	-0.0032	0.0139
Pension or retirement benefits	0.4355	0.4960	0.4428	0.4968	0.0073	0.0176
Dental plan	0.4745	0.4995	0.4733	0.4994	-0.0012	0.0178
Tuition reimbursement	0.2574	0.4374	0.2883	0.4531	0.0309	0.0158
Employment	0.4923	0.5000	0.5242	0.4995	0.0319	0.0112
Sample size	3288		5091		8379	

TABLE 12. Probability of working at amenity-providing firm at week 135 (conditional on employment)

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Health insurance	0.5503	0.4976	0.5754	0.4944	0.0251	0.0155
Paid sick leave	0.4639	0.4988	0.4952	0.5001	0.0313	0.0156
Paid vacation	0.6286	0.4833	0.6341	0.4818	0.0055	0.0147
Childcare assistance	0.1423	0.3495	0.1596	0.3663	0.0173	0.0111
Flexible hours	0.5622	0.4963	0.5839	0.4930	0.0217	0.0153
Employer-provided transportation	0.1841	0.3877	0.1974	0.3981	0.0133	0.0123
Pension or retirement benefits	0.4393	0.4965	0.4754	0.4995	0.0361	0.0156
Dental plan	0.4726	0.4994	0.5001	0.5001	0.0274	0.0156
Tuition reimbursement	0.2676	0.4428	0.3047	0.4604	0.0371	0.0141
Employment	0.5216	0.4996	0.5651	0.4958	0.0435	0.0111
Sample size	3288		5091		8379	

TABLE 13. Probability of working at amenity-providing firm at week 180 (conditional on employment)

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Health insurance	0.5619	0.4963	0.5945	0.4911	0.0326	0.0151
Paid sick leave	0.4752	0.4995	0.5187	0.4997	0.0435	0.0153
Paid vacation	0.6363	0.4812	0.6488	0.4774	0.0126	0.0143
Childcare assistance	0.1441	0.3513	0.1680	0.3739	0.0239	0.0110
Flexible hours	0.6038	0.4892	0.6088	0.4881	0.0049	0.0147
Employer-provided transportation	0.1877	0.3906	0.2038	0.4029	0.0160	0.0121
Pension or retirement benefits	0.4500	0.4976	0.4883	0.5000	0.0382	0.0152
Dental plan	0.4882	0.5000	0.5175	0.4998	0.0293	0.0153
Tuition reimbursement	0.2951	0.4562	0.3150	0.4646	0.0199	0.0140
Employment	0.5510	0.4975	0.5902	0.4918	0.0393	0.0111
Sample size	3288		5091		8379	

TABLE 14. Probability of working at amenity-providing firm at week 208 (conditional on employment)

D.5. Distribution of Job Types. Appendix D.5 presents the distribution of workers by the amenity category their job falls into, as discussed in Footnote 35, for all weeks of interest (90, 135, 180 and 208). The sample size in these distributions decreases to 6,232 individuals (=2,454 control units + 3,778 treated units). Constructing these distributions requires restricting the sample to workers who have non-missing amenity status for all three amenities simultaneously across weeks of interest. This is a stronger restriction than only requiring non-missing amenity status across weeks on a per amenity basis, as is the case for our primary sample of analysis.

	Control		Treated		Difference		Pooled	
	Mean	S.D.	Mean	S.D.	Difference	S.E.	ln(Hourly wage)	S.D.
H=1, R=1, V=1	0.2914	0.4547	0.3352	0.4722	0.0438	0.0198	1.9757	0.3192
H=1, R=1, V=0	0.0303	0.1715	0.0192	0.1375	-0.0110	0.0068	1.9308	0.3114
H=1, R=0, V=1	0.1123	0.3160	0.1073	0.3096	-0.0050	0.0135	1.8983	0.2501
H=1, R=0, V=0	0.0500	0.2180	0.0485	0.2149	-0.0015	0.0093	1.8232	0.2830
H=0, R=1, V=1	0.0288	0.1672	0.0341	0.1817	0.0054	0.0074	1.8834	0.2968
H=0, R=1, V=0	0.0195	0.1383	0.0129	0.1130	-0.0066	0.0055	1.8118	0.2250
H=0, R=0, V=1	0.0758	0.2648	0.0913	0.2881	0.0155	0.0118	1.7606	0.3116
H=0, R=0, V=0	0.3920	0.4885	0.3514	0.4776	-0.0406	0.0208	1.7558	0.3098
Sample size	2454		3778				6232	

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

TABLE 15. Distribution of job types at week 90 (conditional on employment)

	Control		Treated		Difference		Pooled	
	Mean	S.D.	Mean	S.D.	Difference	S.E.	ln(Hourly wage)	S.D.
H=1, R=1, V=1	0.3598	0.4802	0.3806	0.4857	0.0208	0.0192	2.0418	0.3401
H=1, R=1, V=0	0.0262	0.1597	0.0270	0.1620	0.0008	0.0064	2.0489	0.3276
H=1, R=0, V=1	0.1118	0.3152	0.1007	0.3011	-0.0110	0.0123	1.9947	0.3217
H=1, R=0, V=0	0.0345	0.1825	0.0445	0.2063	0.0101	0.0076	1.8928	0.3839
H=0, R=1, V=1	0.0302	0.1713	0.0284	0.1661	-0.0019	0.0067	1.9083	0.2917
H=0, R=1, V=0	0.0222	0.1475	0.0188	0.1357	-0.0035	0.0057	1.8744	0.3267
H=0, R=0, V=1	0.0843	0.2779	0.0625	0.2421	-0.0218	0.0106	1.8479	0.3182
H=0, R=0, V=0	0.3311	0.4709	0.3376	0.4730	0.0064	0.0188	1.8545	0.3716
Sample size	2454		3778				6232	

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

TABLE 16. Distribution of job types at week 135 (conditional on employment)

	Control		Treated		Difference		Pooled	
	Mean	S.D.	Mean	S.D.	Difference	S.E.	ln(Hourly wage)	S.D.
H=1, R=1, V=1	0.3648	0.4816	0.3899	0.4878	0.0251	0.0177	2.1175	0.3441
H=1, R=1, V=0	0.0202	0.1409	0.0324	0.1772	0.0122	0.0057	2.1131	0.3140
H=1, R=0, V=1	0.1208	0.3260	0.1045	0.3059	-0.0163	0.0117	2.0218	0.3153
H=1, R=0, V=0	0.0386	0.1928	0.0381	0.1915	-0.0005	0.0070	1.8936	0.4835
H=0, R=1, V=1	0.0289	0.1677	0.0306	0.1721	0.0016	0.0062	1.9896	0.3512
H=0, R=1, V=0	0.0114	0.1060	0.0183	0.1340	0.0069	0.0043	1.9150	0.3633
H=0, R=0, V=1	0.0851	0.2792	0.0711	0.2571	-0.0140	0.0099	1.8927	0.3296
H=0, R=0, V=0	0.3301	0.4704	0.3151	0.4647	-0.0149	0.0171	1.8606	0.4033
Sample size	2454		3778				6232	

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

TABLE 17. Distribution of job types at week 180 (conditional on employment)

	Control		Treated		Difference		Pooled	
	Mean	S.D.	Mean	S.D.	Difference	S.E.	ln(Hourly wage)	S.D.
H=1, R=1, V=1	0.3663	0.4820	0.4083	0.4916	0.0420	0.0172	2.1271	0.3271
H=1, R=1, V=0	0.0256	0.1579	0.0298	0.1702	0.0043	0.0057	2.0602	0.3268
H=1, R=0, V=1	0.1198	0.3248	0.1073	0.3096	-0.0125	0.0113	2.0046	0.3436
H=1, R=0, V=0	0.0418	0.2003	0.0387	0.1929	-0.0031	0.0070	1.9273	0.4623
H=0, R=1, V=1	0.0291	0.1681	0.0276	0.1639	-0.0015	0.0059	1.9958	0.3719
H=0, R=1, V=0	0.0133	0.1148	0.0170	0.1293	0.0037	0.0043	1.9189	0.3324
H=0, R=0, V=1	0.0875	0.2827	0.0717	0.2580	-0.0158	0.0097	1.9253	0.2689
H=0, R=0, V=0	0.3166	0.4653	0.2996	0.4582	-0.0170	0.0163	1.8768	0.4223
Sample size	2454		3778				6232	

Notes: Hourly wage calculated as weekly earnings divided by weekly hours for the employed.

TABLE 18. Distribution of job types at week 208 (conditional on employment)

D.6. Replication of Lee (2009) Bounds. As discussed in Section 5, Appendix Table 19 reports our replication of the bounds reported in Lee (2009) for weeks 90, 135, 180 and 208 along with the trimming proportion $p \equiv \mathbb{P}(AE)$, e.g., the share of the always-employed among individuals receiving job training. We do not report bounds for week 45 since we discovered that the monotonicity assumption is violated.

Table 19 reports Lee’s bounds when treating $\ln(\text{hourly wage})$ as a continuous variable (as we do throughout the paper). All quantities are very close to the estimates in Lee (2009). There is a small difference that arises in the bounds due to Lee’s use of vingtiles of $\ln(\text{hourly wage})$. Table 20 shows that when we use vingtiles of $\ln(\text{hourly wage})$, the bounds are identical to the ones reported in Lee (2009).

	$\mathbb{P}[D > 0 Z = 0]$	$\mathbb{P}[D > 0 Z = 1]$	p	$\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} D_0 > 0, D_1 > 0]$	
				lower	upper
Week 90	0.4600	0.4601	0.0003	0.0468	0.0484
Week 135	0.5173	0.5451	0.0509	-0.0072	0.0842
Week 180	0.5403	0.5825	0.0724	-0.0325	0.0901
Week 208	0.5655	0.6068	0.0680	-0.0217	0.0989

Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. Propensity scores and trimming proportion are numerically equivalent to Lee; slight numerical difference in bounds occurs as Lee uses vingtiles of $\ln(\text{hourly wage})$ and we do not. See Table 20 for identical treatment bounds to Lee.

TABLE 19. Lee’s bounds: continuous $\ln(\text{hourly wage})$

	$\mathbb{P}[D > 0 Z = 0]$	$\mathbb{P}[D > 0 Z = 1]$	p	$\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} D_0 > 0, D_1 > 0]$	
				lower	upper
Week 90	0.4600	0.4601	0.0003	0.0423	0.0428
Week 135	0.5173	0.5451	0.0509	-0.0159	0.0757
Week 180	0.5403	0.5825	0.0724	-0.0325	0.0868
Week 208	0.5655	0.6068	0.0680	-0.0194	0.0933

Notes: Treatment bounds are for vingtiles of $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed.

TABLE 20. Lee’s bounds: vingtiles of $\ln(\text{hourly wage})$

D.7. Identified Sets for Response Types. Appendix D.7 provides propensity scores for when classifying firm type based on the provision of paid vacation and retirement/pension benefits. Appendix D.7 also provides the identified sets for $p_{L,L}, p_{H,H}$ for all weeks and amenities, with the exception of week 90 for health insurance which is provided in the main text.

	$\mathbb{P}[D = H Z = 0]$	$\mathbb{P}[D = H Z = 1]$	$\mathbb{P}[D = L Z = 0]$	$\mathbb{P}[D = L Z = 1]$
Week 90	0.2470	0.2673	0.2129	0.1928
Week 135	0.3102	0.3209	0.2071	0.2241
Week 180	0.3334	0.3589	0.2069	0.2236
Week 208	0.3516	0.3839	0.2139	0.2229

TABLE 21. Propensity scores by week, amenity: paid vacation

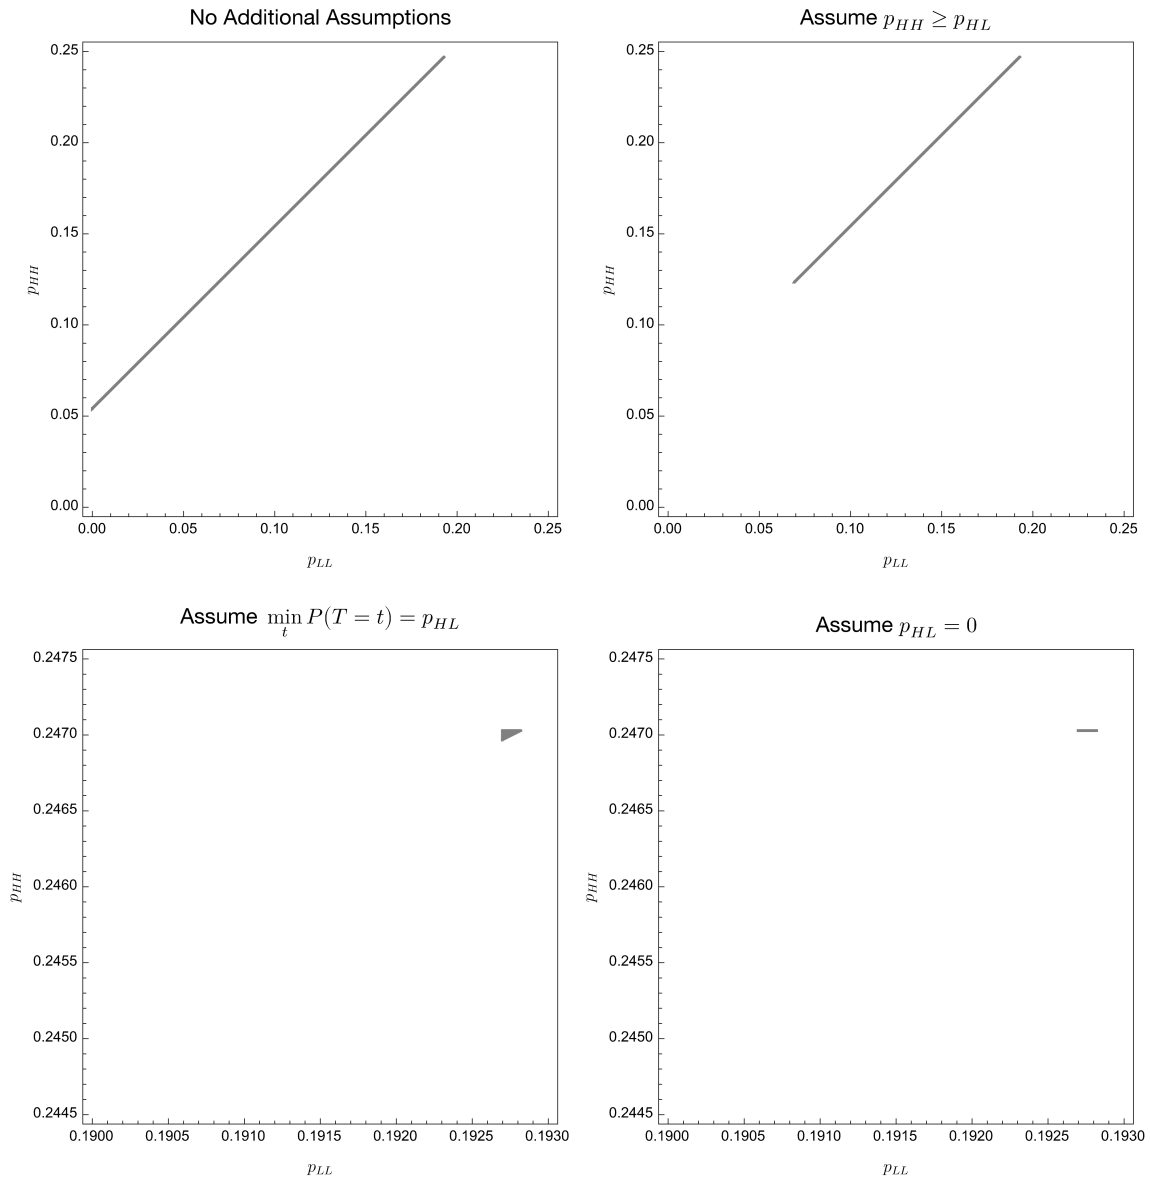


FIGURE 23. Identified set for $p_{L,L}, p_{H,H}$ at week 90. Amenity=paid vacation. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2. The scale of the axis in the bottom two panels is shrunk to see more clearly the identified set.

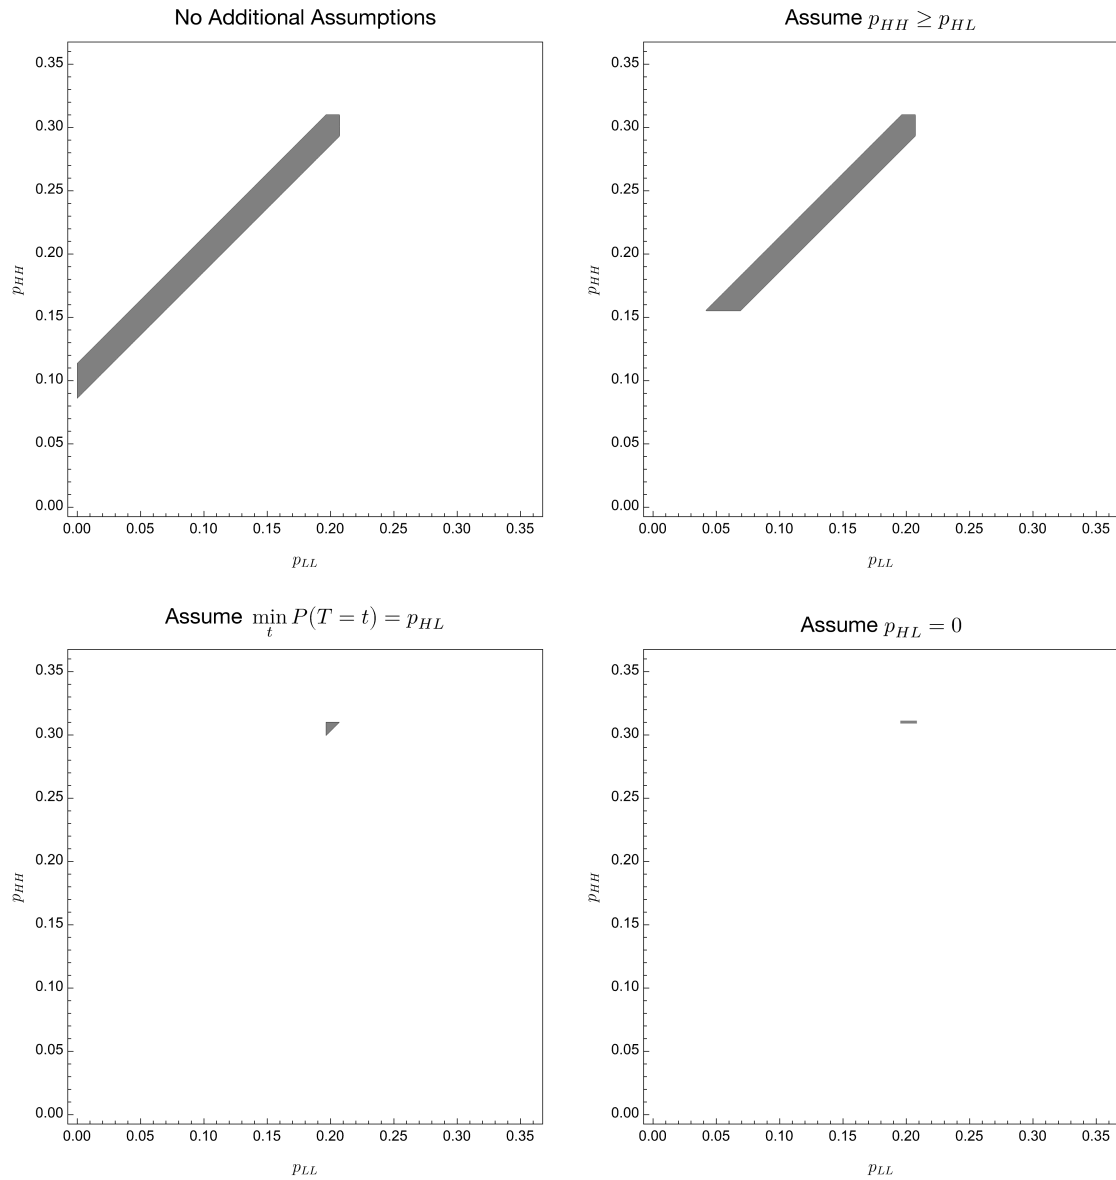


FIGURE 24. Identified set for $p_{L,L}, p_{H,H}$ at week 135. Amenity=paid vacation. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

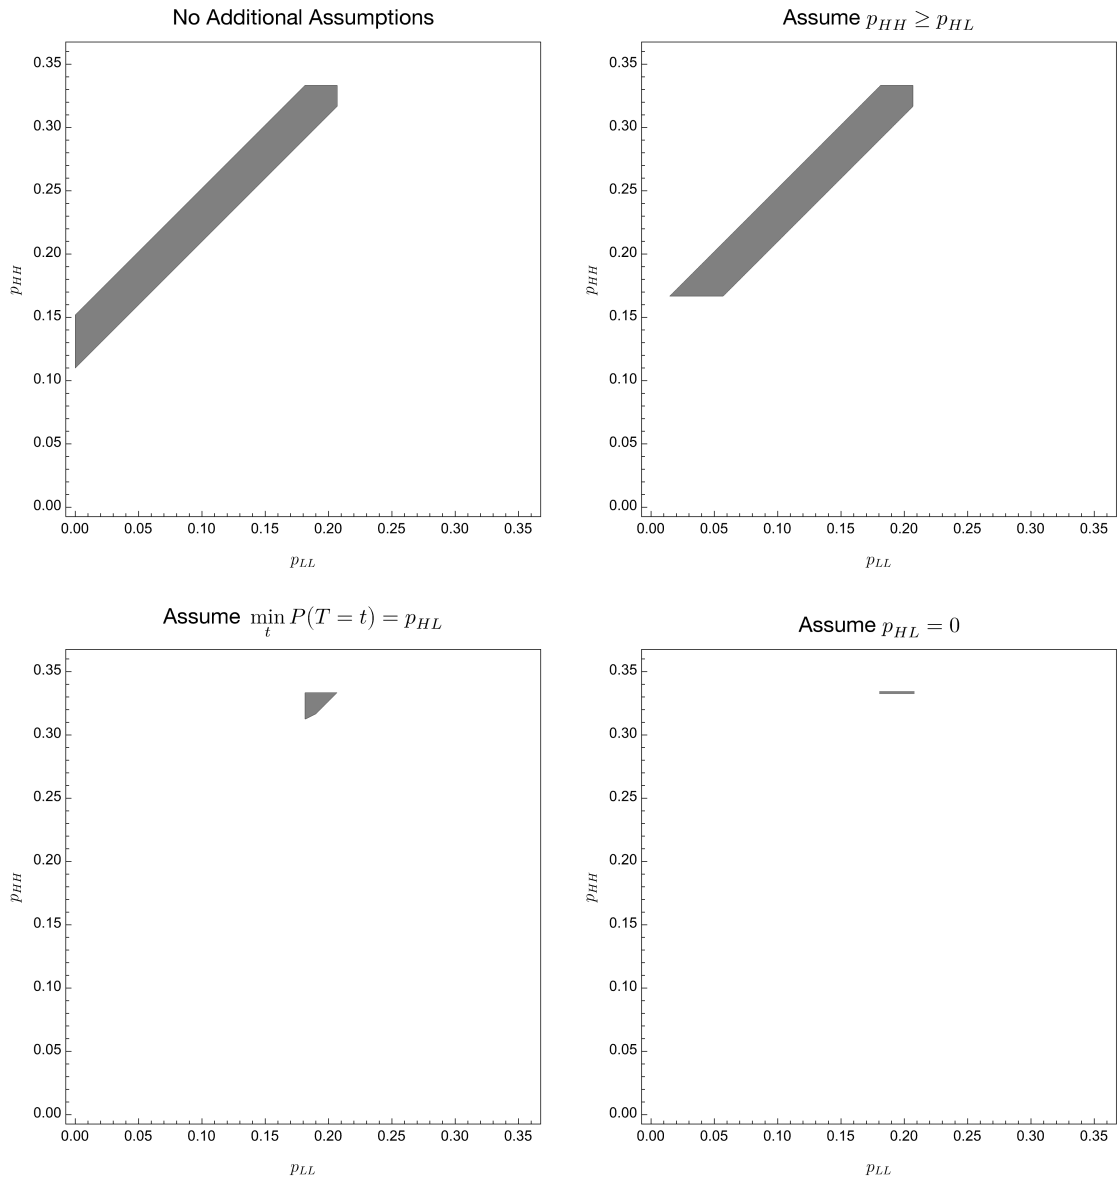


FIGURE 25. Identified set for $p_{L,L}, p_{H,H}$ at week 180. Amenity=paid vacation. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.




FIGURE 26. Identified set for $p_{L,L}, p_{H,H}$ at week 208. Amenity=paid vacation. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

	$\mathbb{P}[D = H Z = 0]$	$\mathbb{P}[D = H Z = 1]$	$\mathbb{P}[D = L Z = 0]$	$\mathbb{P}[D = L Z = 1]$
Week 90	0.1706	0.1852	0.2894	0.2749
Week 135	0.2275	0.2475	0.2898	0.2976
Week 180	0.2309	0.2748	0.3094	0.3076
Week 208	0.2460	0.2930	0.3195	0.3138

TABLE 22. Propensity scores by week, amenity: retirement/pension benefits




FIGURE 27. Identified set for $p_{L,L}, p_{H,H}$ at week 90. Amenity=retirement/pension benefits. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2. The scale of the axis in the bottom two panels is shrunk to see more clearly the identified set.




FIGURE 28. Identified set for $p_{L,L}, p_{H,H}$ at week 135. Amenity=retirement/pension benefits. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

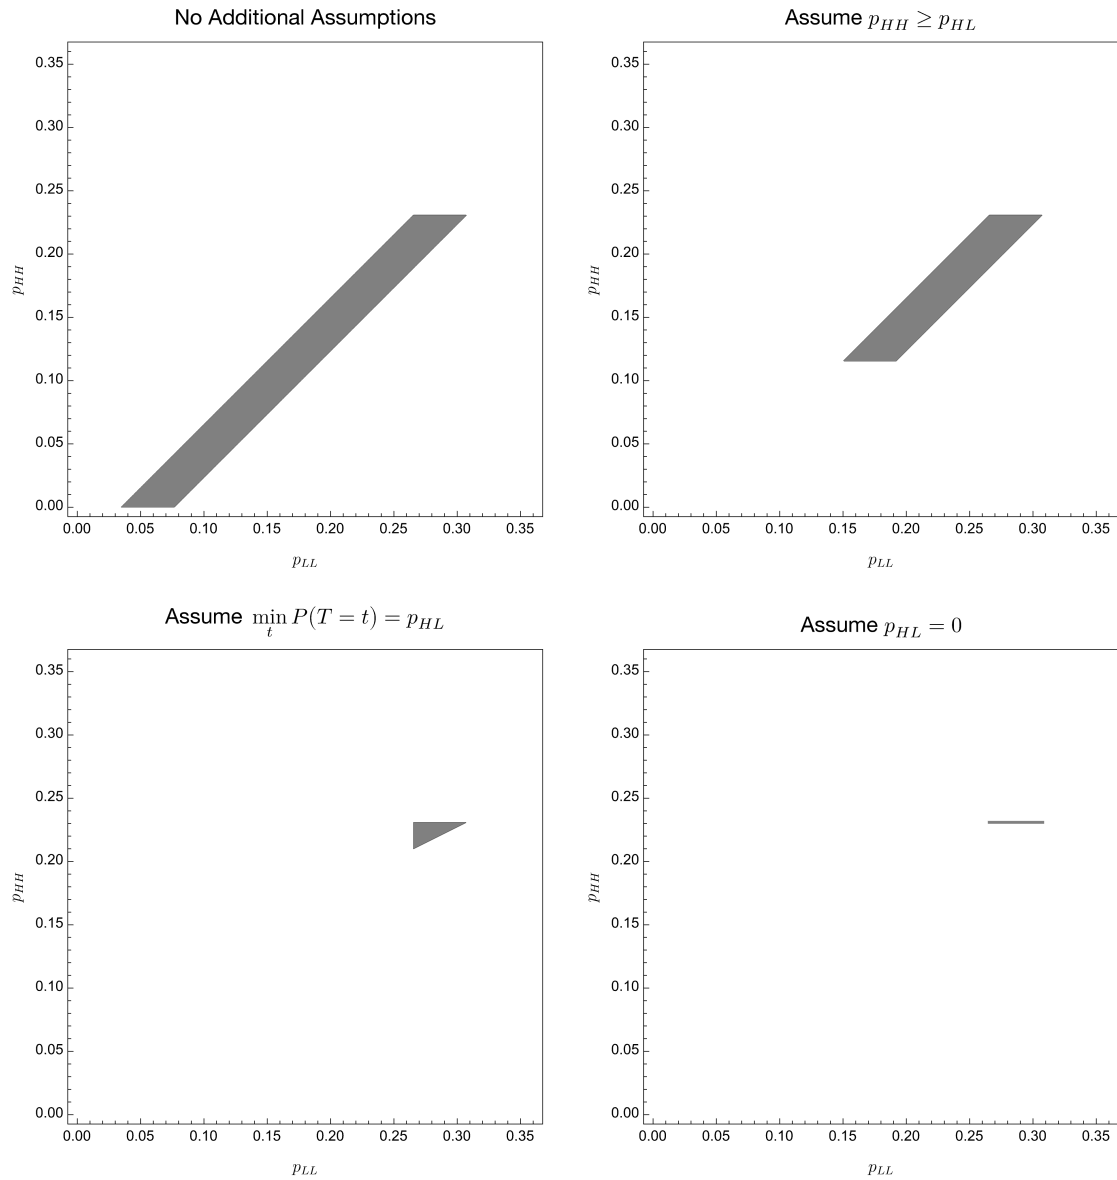


FIGURE 29. Identified set for $p_{L,L}, p_{H,H}$ at week 180. Amenity=retirement/pension benefits. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

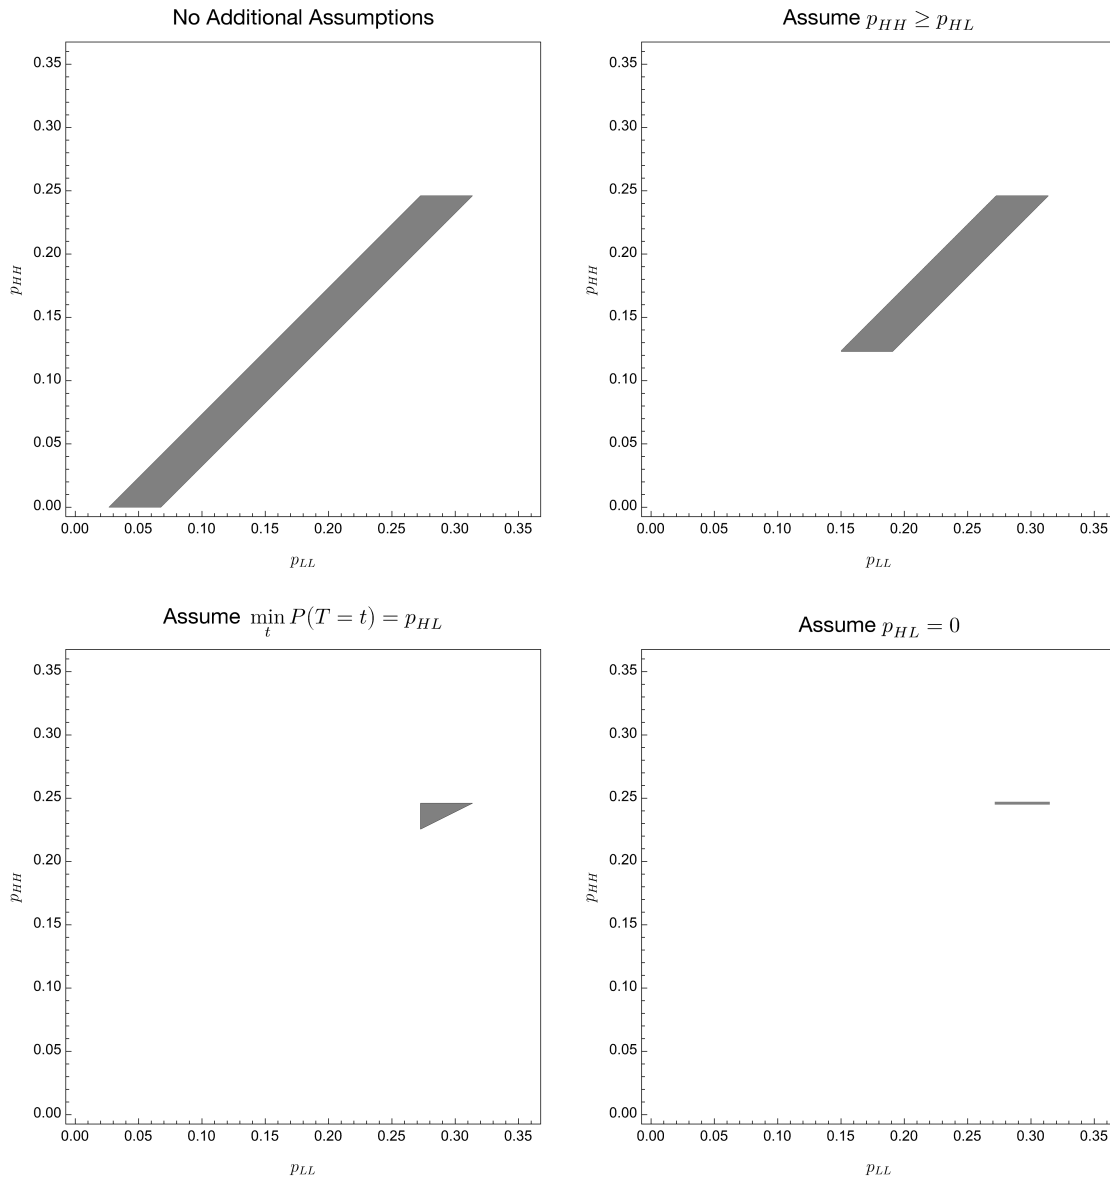


FIGURE 30. Identified set for $p_{L,L}, p_{H,H}$ at week 208. Amenity=retirement/pension benefits. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

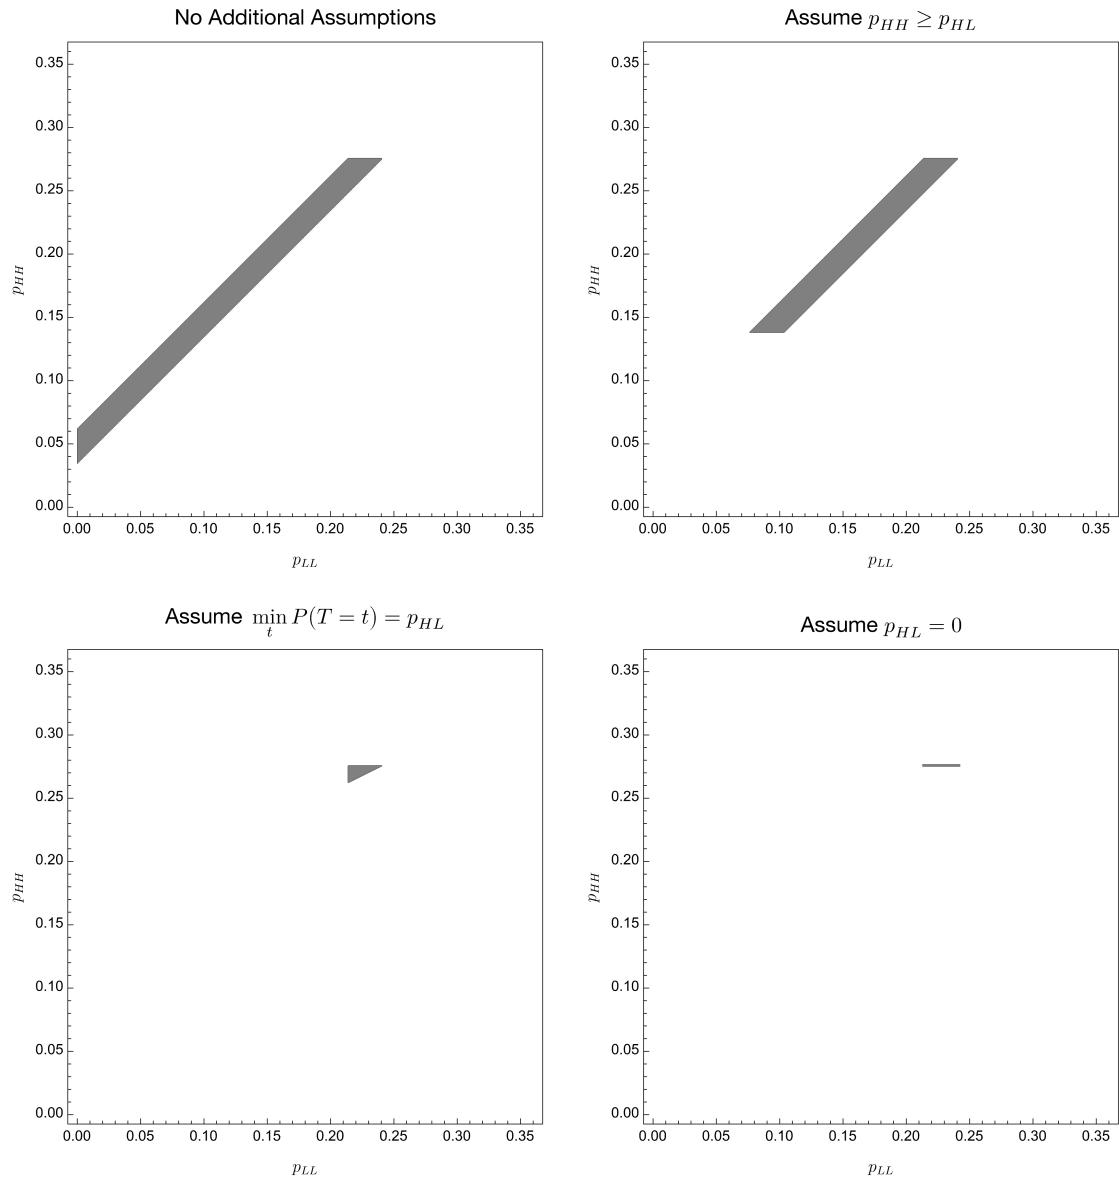


FIGURE 31. Identified set for $p_{L,L}, p_{H,H}$ at week 135. Amenity=health insurance. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.




FIGURE 32. Identified set for $p_{L,L}, p_{H,H}$ at week 180. Amenity=health insurance. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.




FIGURE 33. Identified set for $p_{L,L}, p_{H,H}$ at week 208. Amenity=health insurance. The first panel illustrates the identified set when $\mathcal{R}_T = \{\text{Assumption 2}\}$. The remaining panels illustrate the identified set of additional assumptions imposed over Assumption 2.

D.8. **Multilayered Bounds.** Appendix D.8 provides aggregate multilayered bounds for all classifications of firm type, and within-firm-type multilayered bounds for when classifying firm type based on the provision of paid vacation and retirement/pension benefits.

	p_{HH}^*	p_{LL}^*	p_{HH}^*	p_{LL}^*	$\sum_{d \in \{L, H\}} \frac{p_{d,d}}{p_{HH} + p_{LL}} \mathbb{E}[Y_{1,d} - Y_{0,d} T = (d, d)]$	
Week 90	(for lower bound)		(for upper bound)		lower	upper
Baseline	0.0014	0.0004	0.0014	0.0004	-2.4634	2.4588
$p_{H,H} \geq p_{H,L}$	0.1120	0.1109	0.1120	0.1109	-0.4106	0.4780
(H,L) is smallest type	0.2235	0.2224	0.2235	0.2224	-0.0136	0.0749
$p_{H,L} = 0$	0.2235	0.2224	0.2235	0.2224	-0.0136	0.0749
Week 135						
Baseline	0.0366	0.0022	0.0392	0.0048	-1.1602	1.2762
$p_{H,H} \geq p_{H,L}$	0.1380	0.0758	0.1380	0.0758	-0.5423	0.6017
(H,L) is smallest type	0.2614	0.2127	0.2614	0.2127	-0.1179	0.1532
$p_{H,L} = 0$	0.2754	0.2132	0.2754	0.2132	-0.0833	0.1165
Week 180						
Baseline	0.0475	0.0046	0.0485	0.0056	-1.1369	1.2631
$p_{H,H} \geq p_{H,L}$	0.1471	0.0620	0.1471	0.0620	-0.5846	0.6407
(H,L) is smallest type	0.2727	0.2081	0.2727	0.2081	-0.1607	0.1900
$p_{H,L} = 0$	0.2937	0.2086	0.2937	0.2086	-0.1101	0.1397
Week 208						
Baseline	0.0705	0.0072	0.0671	0.0038	-0.9987	1.0704
$p_{H,H} \geq p_{H,L}$	0.1571	0.0525	0.1571	0.0525	-0.5932	0.6436
(H,L) is smallest type	0.2928	0.2087	0.2928	0.2087	-0.1532	0.1941
$p_{H,L} = 0$	0.3138	0.2092	0.3138	0.2092	-0.1049	0.1461

Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_t^* is the optimal value of p_t over the joint identified set for response-types under the given assumption; p_t^* may differ for lower and upper bounds and is therefore presented separately.

TABLE 23. Aggregate multilayered bounds for Job Corps; health insurance amenity.

Week 90	p_{HH}^*	p_{LL}^*	$\mathbb{E}(Y_{1,H} - Y_{0,H} T = (H, H))$		$\mathbb{E}(Y_{1,L} - Y_{0,L} T = (L, L))$	
			lower	upper	lower	upper
Baseline	0.0542	0.0000	-0.8794	0.9952		
$p_{H,H} \geq p_{H,L}$	0.1235	0.0692	-0.5033	0.5722	-0.5478	0.6430
(H,L) is smallest type	0.2470	0.1927	-0.0228	0.1002	-0.0242	0.1061
$p_{H,L} = 0$	0.2470	0.1927	-0.0218	0.0995	-0.0242	0.1061
Week 135						
Baseline	0.0860	0.0000	-0.7956	0.8818		
$p_{H,H} \geq p_{H,L}$	0.1551	0.0413	-0.5108	0.5835	-0.8611	0.9932
(H,L) is smallest type	0.2994	0.1964	-0.0655	0.1239	-0.1186	0.1420
$p_{H,L} = 0$	0.3102	0.1964	0.0012	0.0725	-0.1186	0.1420
Week 180						
Baseline	0.1099	0.0000	-0.7394	0.7904		
$p_{H,H} \geq p_{H,L}$	0.1667	0.0147	-0.5265	0.5794	-1.6717	1.7737
(H,L) is smallest type	0.3123	0.1814	-0.1244	0.1629	-0.1909	0.2229
$p_{H,L} = 0$	0.3334	0.1814	-0.0361	0.0806	-0.1909	0.2229
Week 208						
Baseline	0.1287	0.0000	-0.6719	0.7683		
$p_{H,H} \geq p_{H,L}$	0.1758	0.0058	-0.5136	0.6070	-2.7579	2.6008
(H,L) is smallest type	0.3310	0.1816	-0.1124	0.1847	-0.2046	0.2294
$p_{H,L} = 0$	0.3516	0.1816	-0.0206	0.1083	-0.2046	0.2294

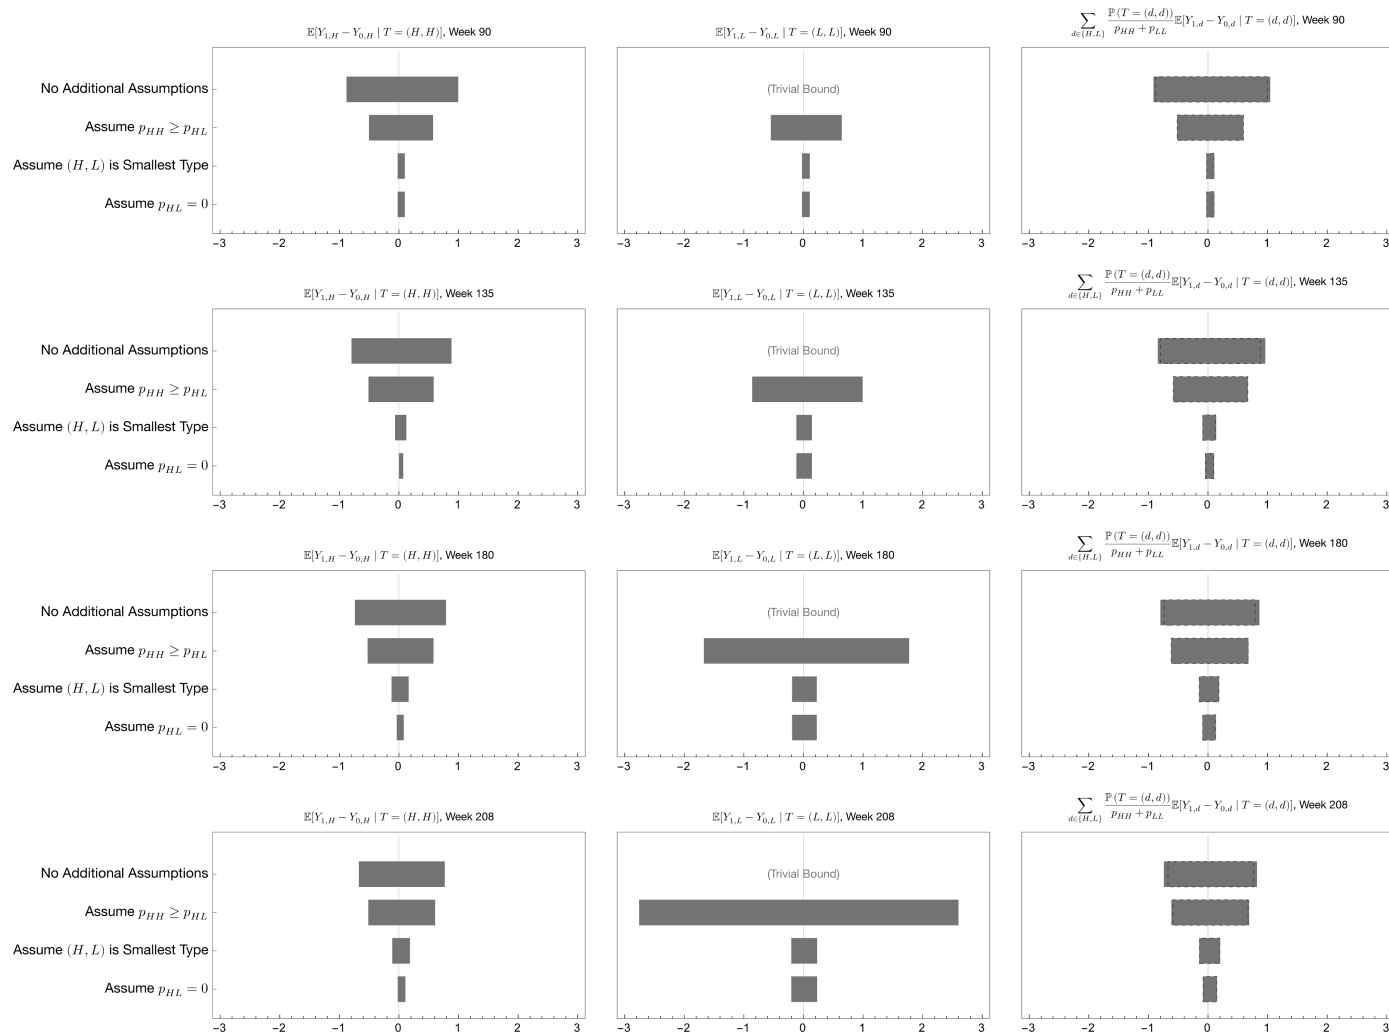
Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_t^* is the minimum value of p_t over the identified set for response-types under the given assumption.

TABLE 24. Multilayered bounds by week, amenity: paid vacation.

	Lower		Upper		$\frac{\sum_{d \in \{L,H\}} p_{d,d}}{\sum_{d' \in \{L,H\}} p_{d',d'}} \mathbb{E}[Y_{1,d} - Y_{0,d} T = (d,d)]$	
	p_{HH}^*	p_{LL}^*	p_{HH}^*	p_{LL}^*	lower	upper
Week 90						
Baseline	0.0577	0.0035	0.0585	0.0043	-0.9155	1.0441
$p_{H,H} \geq p_{H,L}$	0.1235	0.0692	0.1235	0.0692	-0.5192	0.5976
(H,L) is smallest type	0.2466	0.1922	0.2466	0.1922	-0.0284	0.1068
$p_{H,L} = 0$	0.2466	0.1922	0.2466	0.1922	-0.0284	0.1068
Week 135						
Baseline	0.0921	0.0060	0.0933	0.0072	-0.8424	0.9616
$p_{H,H} \geq p_{H,L}$	0.1551	0.0413	0.1551	0.0413	-0.5844	0.6696
(H,L) is smallest type	0.2978	0.1955	0.2978	0.1955	-0.0919	0.1362
$p_{H,L} = 0$	0.3098	0.1960	0.3098	0.1960	-0.0484	0.1022
Week 180						
Baseline	0.1173	0.0075	0.1162	0.0064	-0.7990	0.8620
$p_{H,H} \geq p_{H,L}$	0.1670	0.0228	0.1670	0.0204	-0.6207	0.6776
(H,L) is smallest type	0.3120	0.1805	0.3120	0.1805	-0.1506	0.1868
$p_{H,L} = 0$	0.3330	0.1810	0.3330	0.1810	-0.0935	0.1328
Week 208						
Baseline	0.1374	0.0087	0.1352	0.0065	-0.7404	0.8219
$p_{H,H} \geq p_{H,L}$	0.1761	0.0249	0.1761	0.0186	-0.6148	0.6902
(H,L) is smallest type	0.3302	0.1807	0.3302	0.1807	-0.1477	0.2034
$p_{H,L} = 0$	0.3512	0.1812	0.3512	0.1812	-0.0860	0.1517

Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_i^* is the optimal value of p_i over the joint identified set for response-types under the given assumption; p_i^* may differ for lower and upper bounds and is therefore presented separately.

TABLE 25. Aggregate multilayered bounds by week. Amenity=paid vacation.



Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 34. Multilayered bounds by week. Amenity=paid vacation.

	p_{HH}^*	p_{LL}^*	$\mathbb{E}(Y_{1,H} - Y_{0,H} T = (H, H))$		$\mathbb{E}(Y_{1,L} - Y_{0,L} T = (L, L))$	
			lower	upper	lower	upper
Week 90						
Baseline	0.0000	0.1042			-0.5064	0.5623
$p_{H,H} \geq p_{H,L}$	0.0853	0.1894	-0.4465	0.5337	-0.2568	0.3074
(H,L) is smallest type	0.1705	0.2747	-0.0141	0.0921	-0.0161	0.0628
$p_{H,L} = 0$	0.1706	0.2747	-0.0135	0.0916	-0.0161	0.0628
Week 135						
Baseline	0.0000	0.0423			-0.9636	1.0670
$p_{H,H} \geq p_{H,L}$	0.1138	0.1561	-0.4749	0.5746	-0.4374	0.4505
(H,L) is smallest type	0.2136	0.2699	-0.0854	0.1729	-0.1203	0.1144
$p_{H,L} = 0$	0.2275	0.2699	-0.0143	0.0977	-0.1203	0.1144
Week 180						
Baseline	0.0000	0.0345			-1.1762	1.2907
$p_{H,H} \geq p_{H,L}$	0.1155	0.1500	-0.5247	0.5641	-0.4911	0.5305
(H,L) is smallest type	0.2098	0.2655	-0.1802	0.2097	-0.1655	0.1947
$p_{H,L} = 0$	0.2309	0.2655	-0.0856	0.1062	-0.1655	0.1947
Week 208						
Baseline	0.0000	0.0265			-1.4213	1.4398
$p_{H,H} \geq p_{H,L}$	0.1230	0.1495	-0.5019	0.5491	-0.5103	0.5400
(H,L) is smallest type	0.2254	0.2725	-0.1615	0.2023	-0.1640	0.1975
$p_{H,L} = 0$	0.2460	0.2725	-0.0784	0.1084	-0.1640	0.1975

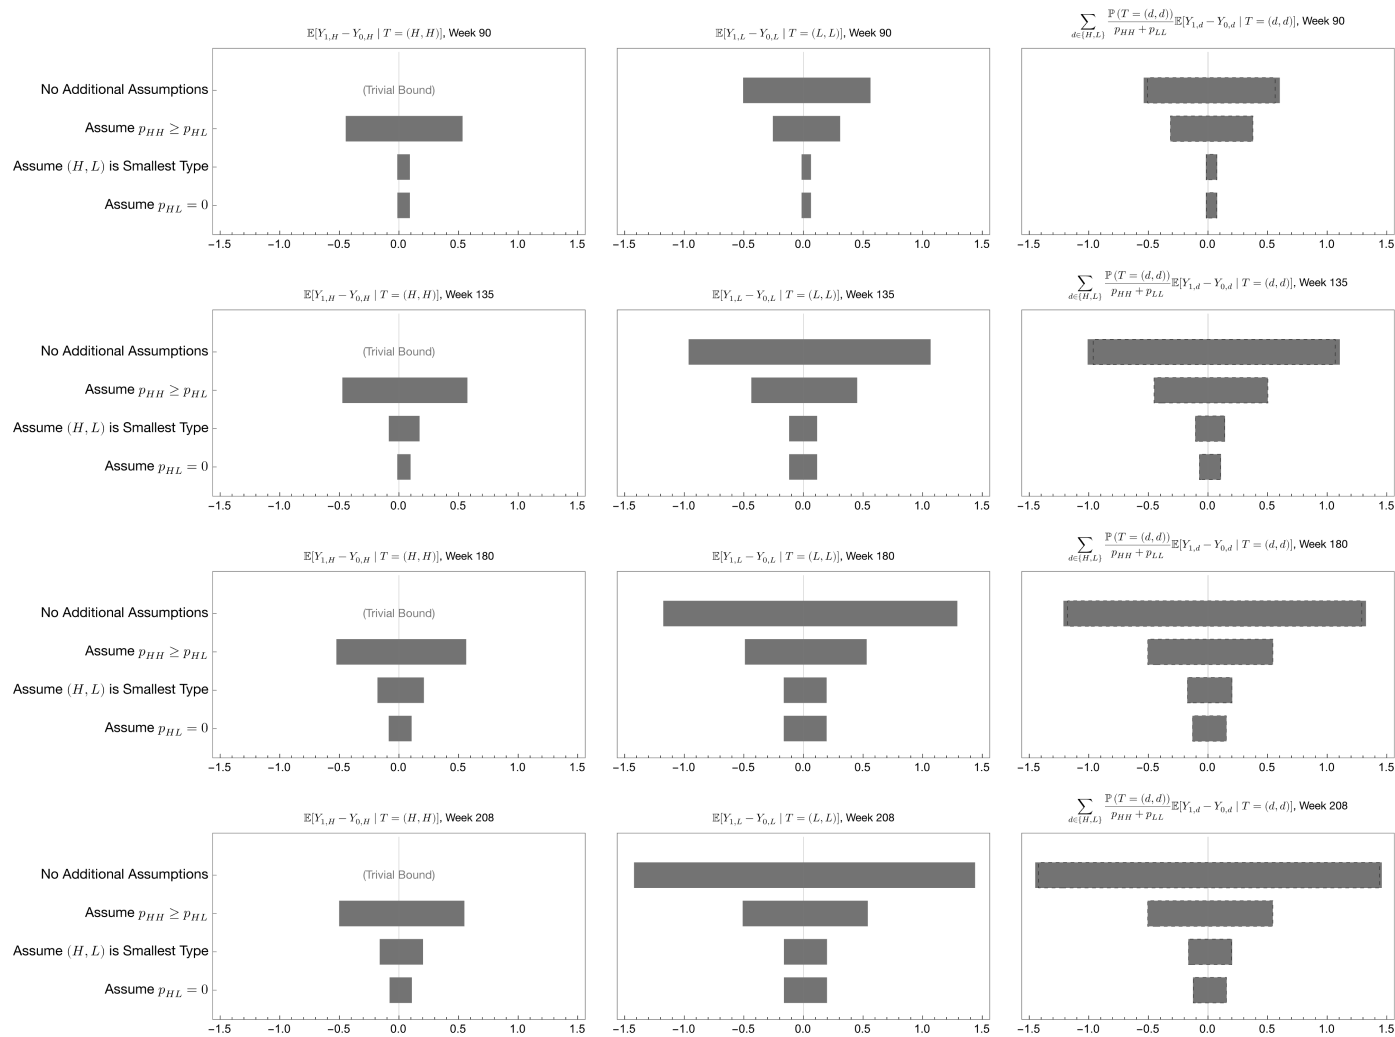
Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_t^* is the minimum value of p_t over the identified set for response-types under the given assumption.

TABLE 26. Multilayered bounds by week, amenity: retirement/pension benefits.

	Lower		Upper		$\frac{\sum_{d \in \{L,H\}} p_{d,d}}{\sum_{d' \in \{L,H\}} p_{d',d'}} \mathbb{E}[Y_{1,d} - Y_{0,d} T = (d,d)]$	
	p_{HH}^*	p_{LL}^*	p_{HH}^*	p_{LL}^*	lower	upper
Week 90						
Baseline	0.0075	0.1117	0.0084	0.1126	-0.5397	0.6025
$p_{H,H} \geq p_{H,L}$	0.0852	0.1894	0.0852	0.1894	-0.3159	0.3779
(H,L) is smallest type	0.1702	0.2744	0.1702	0.2744	-0.0181	0.0773
$p_{H,L} = 0$	0.1702	0.2744	0.1702	0.2744	-0.0181	0.0773
Week 135						
Baseline	0.0032	0.0455	0.0038	0.0461	-1.0101	1.1059
$p_{H,H} \geq p_{H,L}$	0.1141	0.1565	0.1141	0.1565	-0.4519	0.5015
(H,L) is smallest type	0.2131	0.2689	0.2131	0.2689	-0.1074	0.1430
$p_{H,L} = 0$	0.2271	0.2694	0.2271	0.2694	-0.0739	0.1091
Week 180						
Baseline	0.0024	0.0371	0.0024	0.0371	-1.2130	1.3253
$p_{H,H} \geq p_{H,L}$	0.1154	0.1501	0.1154	0.1501	-0.5056	0.5450
(H,L) is smallest type	0.2095	0.2646	0.2095	0.2646	-0.1739	0.2032
$p_{H,L} = 0$	0.2305	0.2651	0.2305	0.2651	-0.1309	0.1555
Week 208						
Baseline	0.0010	0.0276	0.0014	0.0280	-1.4492	1.4578
$p_{H,H} \geq p_{H,L}$	0.1230	0.1496	0.1230	0.1496	-0.5063	0.5439
(H,L) is smallest type	0.2246	0.2716	0.2246	0.2716	-0.1655	0.2024
$p_{H,L} = 0$	0.2456	0.2721	0.2456	0.2721	-0.1257	0.1572

Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed. p_i^* is the optimal value of p_i over the joint identified set for response-types under the given assumption; p_i^* may differ for lower and upper bounds and is therefore presented separately.

TABLE 27. Aggregate multilayered bounds by week. Amenity=retirement/pension benefits.



Notes: Treatment bounds are for $\ln(\text{hourly wage})$; hourly wage calculated as weekly earnings divided by weekly hours for the employed.

FIGURE 35. Multilayered bounds by week. Amenity=pension benefits.

APPENDIX E. ADDITIONAL EMPIRICAL TABLES AND FIGURES FOR WORKADVANCE
RCTs

E.1. Summary Statistics and Differential Sorting of Treatment and Control Workers. This section provides descriptive statistics on the WorkAdvance RCT.

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Female	0.14	0.34	0.18	0.39	0.05	0.03
Adults ≤ 24	0.22	0.42	0.23	0.42	0.01	0.03
Black	0.35	0.48	0.39	0.49	0.03	0.04
HS/GED or less	0.40	0.49	0.43	0.50	0.03	0.04
At baseline:						
Have job	0.28	0.45	0.26	0.44	-0.02	0.03
Quarterly earnings	2257.07	2820.54	1944.51	2495.57	-312.56	201.91
8 quarters post randomization:						
Employment	0.66	0.48	0.67	0.47	0.01	0.04
Share employment in target sector	0.31	0.46	0.44	0.50	0.14	0.04
Quarterly earnings	3703.87	4169.10	4049.91	4003.27	346.04	309.72
Sample size	344		353		697	

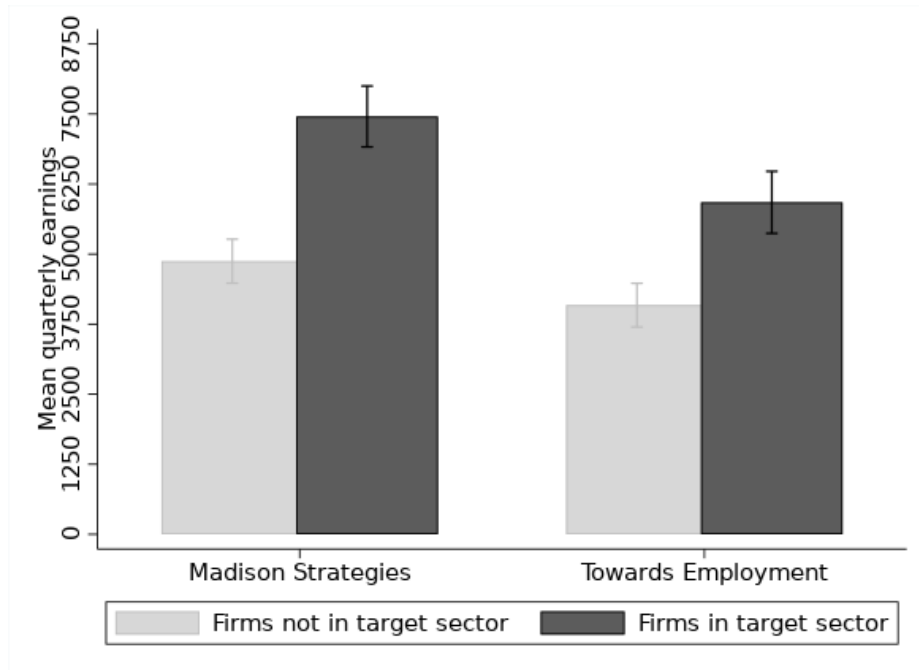
Notes: Quarterly earnings are not conditional on employment (i.e., includes 0s for the unemployed).

TABLE 28. Summary statistics by treatment status, RCT = Madison Strategies

	Control		Treated		Difference	
	Mean	S.D.	Mean	S.D.	Difference	S.E.
Female	0.59	0.49	0.58	0.49	-0.01	0.04
Adults ≤ 24	0.23	0.42	0.22	0.42	-0.01	0.03
Black	0.73	0.45	0.77	0.42	0.05	0.03
HS/GED or less	0.42	0.49	0.44	0.50	0.03	0.04
At baseline:						
Have job	0.27	0.45	0.26	0.44	-0.01	0.03
Quarterly earnings	1504.59	2201.72	1723.72	2452.60	219.12	176.42
8 quarters post randomization:						
Employment	0.62	0.49	0.69	0.46	0.08	0.04
Share employment in target sector	0.47	0.50	0.51	0.50	0.03	0.05
Quarterly earnings	3040.57	4024.21	3497.54	3684.01	456.97	292.04
Sample size	349		349		698	


Notes: Quarterly earnings are not conditional on employment (i.e., includes 0s for the unemployed).

TABLE 29. Summary statistics by treatment status, RCT = Towards Employment



Notes: Quarterly earnings means are among the employed (i.e., do not include 0s for the unemployed).

FIGURE 36. Mean quarterly earnings by firm sector




(a) control units

(b) treated units

Notes: Quarterly earnings are for the employed (i.e., does not include 0s for the unemployed).

FIGURE 37. Cumulative distribution function by firm sector, RCT= Madison Strategies



(a) control units

(b) treated units

Notes: Quarterly earnings are for the employed (i.e., does not include 0s for the unemployed).

FIGURE 38. Cumulative distribution function by firm sector, RCT= Towards Employment

E.2. Aggregate multilayered bounds for WorkAdvance RCT. This section presents aggregate multilayered bounds for the WorkAdvance RCT.

TABLE 30. Aggregate multilayered bounds for WorkAdvance RCT.

Madison Strategies	p_{HH}^* (for lower bound)	p_{LL}^*	p_{HH}^* (for upper bound)	p_{LL}^*	$\sum_{d \in \{L, H\}} \frac{p_{d,d}}{p_{HH} + p_{LL}} \mathbb{E}[Y_{1,d} - Y_{0,d} T = (d, d)]$	
					lower	upper
Baseline	0.0145	0.1710	0.0090	0.1655	-6498.0976	7255.9812
$p_{H,H} \geq p_{H,L}$	0.1010	0.2570	0.1010	0.2570	-4514.3517	4886.1844
(H,L) is smallest type	0.1960	0.3545	0.1960	0.3545	-1593.0171	1848.4113
$p_{H,L} = 0$	0.2010	0.3570	0.2010	0.3570	-1459.5360	1688.9594
Towards Employment						
Baseline	0.0000	0.0000	0.0000	0.0000		
$p_{H,H} \geq p_{H,L}$	0.1450	0.1170	0.1450	0.1170	-6142.9166	6288.9131
(H,L) is smallest type	0.2500	0.2590	0.2500	0.2590	-2318.5475	2613.7796
$p_{H,L} = 0$	0.2900	0.2620	0.2900	0.2620	-1562.2141	1840.9057

Notes: Outcome is quarterly wages. p_i^* is the minimum value of p_i over the identified set for response-types under the given assumption; p_i^* may differ for lower and upper bounds and is therefore presented separately.

E.3. Lee (2009) Bounds. This section generates Lee’s bounds for the WorkAdvance RCT.

	$\mathbb{P}[D > 0 Z = 0]$	$\mathbb{P}[D > 0 Z = 1]$	p	$\mathbb{E}[Y_{1,D_1} - Y_{0,D_0} D_0 > 0, D_1 > 0]$	
				lower	upper
Madison Strategies	0.6570	0.6686	0.0173	244.1928	524.9793
Towards Employment	0.6160	0.6934	0.1116	-676.8651	705.6605

Notes: Treatment bounds are for quarterly wages.

TABLE 31. Lee’s bounds for WorkAdvance RCT.

APPENDIX F. ‘TOP 5’ PAPERS (POTENTIALLY) COLLAPSING MULTILAYERED SELECTION
TO SINGLE LAYERED SELECTION

In a literature survey which we detail in Section 1, we counted 56 papers published in ‘top 5’ general interest economic journals that cited Lee (2009) and 42 that empirically implemented Lee bounds. Table 32 details 7 of these papers that feature multilayered selection, where researchers simplified the sample selection problem by collapsing it to a single dimension.³⁶

Paper			Sample Selection	
Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Daruich et al. (2023)</i>				
Study 2001 Italian reform lifting constraints on the employment of temporary contract workers but maintaining employment protection laws for permanent contract employees; exploit staggered implementation across collective bargaining agreements.	See “Application”.	Individuals’ earnings.	Only observe employment outcomes after labor market entry; utilize Lee bounds to address that the reform affects the entry margin into the labor market.	Conditional on labor market entry, altering temporary contract worker regulations likely also affects worker sorting across industries and/or firms; indeed this paper finds meaningful changes in the shares of temporary contract workers in certain industries.

³⁶Only outcomes for which Lee’s bounds are estimated are noted in column 3.

... continued

Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Cullen and Pakzad-Hurson (2023)</i>				
Study state-level laws in the U.S. protecting the right of private sector workers to discuss salary information with co-workers.	See “Application”.	Worker wages.	Treatment effects could be driven by compositional changes of private sector workers if high-paid (low-paid) workers disproportionately leave (join) the private sector; estimate Lee bounds to address this challenge.	Conditional on private sector entry, treatment may also affect the sorting of workers across industries or firms within the private sector; for example, knowledge of co-workers’ salaries could cause workers to sort to firms with flatter pay hierarchies (i.e., fairness concerns).

... continued

Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Bianchi and Giorcelli (2022)</i>				
Study long-term and spillover effects of management interventions on firm performance; estimate effect of the Training Within Industry program, a U.S. government training program intended to be provided to all firms involved in war production between 1940 and 1945.	Exploit that constraints resulted in only 7% of applicant firms receiving full training, 48% receiving no training and the remainder receiving partial training; this paper compares applicant firms who received training to applicant firms who did not.	Firm total factor productivity.	Estimate Lee bounds to address the higher attrition rate of untrained firms; treated firms had 90% survival rate at least 10 years following treatment whereas control firms only had 64% survival rate.	Conditional on firm survival, training may also affect the sorting of firms across industries or other important dimensions. This is particularly plausible in this paper's setting where: (i) trained firms undertook structural changes transforming them into larger and more complex organizations; (ii) this paper estimates treatment bounds for outcomes observed post-Second World War, when many firms would have plausibly switched industries (i.e., left war production).

... continued

Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Fink et al. (2020)</i>				
Complete experiment offering subsidized loans to randomly selected villages in rural Zambia where farmers suffer from liquidity constraints in the months prior to harvest (lean season).	Offered cash and maize loans at the start of the lean season (January) with repayment due at harvest (July) in either cash, maize or both.	Individual- and village-level earnings.	Estimate Lee bounds to address that the likelihood of entering the labor market decreases with the loan treatment.	Conditional on entry to the labor market, treatment has potential to affect the types of jobs individuals accept; this is particularly plausible in this paper's setting as labor sales occur within villages between better- and worse-off farmers at individually negotiated rates.

... continued

Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Giorcelli (2019)</i>				
Study long-run effects of management on firm performance; estimate effects of US Technical Assistance and Productivity Program (USTAPP) which provided management training and technologically advanced machines to Italian firms from 1952 to 1958.	Exploit unexpected USTAPP budget cut which occurred after the firm application submission and review period which resulted in many firms not receiving training; compare applicant firms who received training to those who did not.	Firm-level: sales, number of employees, total factor productivity revenue.	Estimate Lee bounds to address the treatment-control difference in firm survival probability.	As in Bianchi and Giorcelli (2022), conditional on firm survival, management training likely affects the sorting of firms across industries and/or other important dimensions.

... continued

Application	Treatment	Outcomes	Layer addressed	Additional layer
<i>Fisman et al. (2017)</i>				
Estimate effect of cultural proximity (e.g., shared codes, language, religion) on loan outcomes for lenders and borrowers using dyadic data on religion and caste for lending officers and borrowers from a state-owned Indian bank.	Exploit lending officer rotation policy providing variation in officer-borrower matching.	Cultural group-level: amount of debt received, total number of borrowers and average loan size.	Estimate Lee bounds as column 3 outcomes are only observed conditional on a group receiving credit.	Conditional on a group receiving credit, “same group matches” also plausibly affect the type of loans a group receive. For example, “same group match” borrowers may receive favorable loan terms; this paper indeed notes the potential for these effects but is constrained by data limitations.
<i>Blanco et al. (2013)</i>				
Estimate bounds on average and quantile treatment effects of Job Corps.	Job Corps program; see description in Section 5.1.	Worker wages.	Estimate Lee bounds to address that Job Corps training may affect labor supply along the extensive margin.	See Appendix D.3.

TABLE 32. ‘Top 5’ papers potentially collapsing multilayered sample selection to a single dimension