

NBER WORKING PAPER SERIES

DO WORKFORCE DEVELOPMENT PROGRAMS BRIDGE THE SKILLS GAP?

Eleanor W. Dillon
Lisa B. Kahn
Joanna Venator
Michael Dalton

Working Paper 34012
<http://www.nber.org/papers/w34012>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2025, Revised April 2026

We thank seminar participants at Arizona State University, Carnegie Mellon, the Cleveland Fed, Columbia, the Family and Education Workshop 2025, the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Dallas, the Markets with Frictions Tinos Edition, NBER Organizational Economics, SITE: The Micro and Macro of Labor Markets, Society of Labor Economics, San Diego State University, SEMML, University of Michigan, and the University of Rochester. We are indebted to Kalab Begna, Edoardo Bella, Klint Mane, and Sid Satya for excellent research assistance and to the broader team of research assistants who contributed to data collection and validation. This project was funded in part by generous support from the Smith Richardson Foundation. All views expressed in this paper are those of the authors and do not necessarily reflect the views or policies of the U.S. Bureau of Labor Statistics or 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Eleanor W. Dillon, Lisa B. Kahn, Joanna Venator, and Michael Dalton. 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.

Do Workforce Development Programs Bridge the Skills Gap?
Eleanor W. Dillon, Lisa B. Kahn, Joanna Venator, and Michael Dalton
NBER Working Paper No. 34012
July 2025, Revised April 2026
JEL No. J24, J42, J68, M53

ABSTRACT

Most U.S. states have workforce development programs that offer firms grants to train their own workers. We create unique data linkages between participating firms, employment, and vacancies to explore the determinants and consequences of such programs. Training grants are more prevalent in markets where firms face greater employee poaching risk, suggesting these programs help overcome a market failure in updating worker skills. After training, firms experience prolonged employment growth and downskilling in their job posts, relative to a matched control group. Training appears to help firms move toward optimal scale and expand opportunities for less skilled workers.

Eleanor W. Dillon
Microsoft Research
eldillon@microsoft.com

Lisa B. Kahn
University of Rochester
Department of Economics
and NBER
lisa.kahn@rochester.edu

Joanna Venator
Boston College
Department of Economics
venator@bc.edu

Michael Dalton
Bureau of Labor Statistics
Dalton.Michael@bls.gov

1 Introduction

Technology and international trade have changed the nature of work in the United States, shifting demand towards workers with a college degree and compressing the bottom of the earnings distribution. At the same time, employers commonly lament a “skills shortage,” a problem that was exacerbated in the extremely tight labor markets of the COVID recovery.¹ Formal schooling is a proposed front-line solution for these problems, as differences in employment and wages between those with and without a college degree are stark (Abraham and Kearney, 2020; Autor, 2019). Nonetheless, educational attainment has stagnated, and, due to the pace of technological change, many workers will need to acquire new skills throughout the course of their careers (Murphy and Topel, 2016; Goldin and Katz, 2008). Unfortunately, public-sector job training programs have historically had, at best, mixed success at offering an alternative to formal schooling.

Private-sector training programs may be more effective because employers know best which skills they need. However, employers will be reluctant to pay to train workers in general skills for fear that their investment will be poached away (Becker, 1964), and workers may not have the resources or knowledge to invest in training themselves (Caliendo et al., 2022). Public-private partnerships, characterized by employer-driven training funded at least in part by the public sector, may help to overcome these frictions. Federal funding for these partnerships has increased in the last decade and most states in the U.S. have at least one program whereby employers apply for grants funded by the government to train their incumbent (either longstanding or newly hired) workers. Despite this growing interest from policymakers, there are no broad-based studies about how and why these programs operate and whether they are successful.

What can the presence and effect of public-private incumbent worker training programs tell us about frictions in worker training and skills gaps? In this paper, we assemble a new dataset of participating firms linked to two rich firm-level datasets – the Quarterly Census of Employment and Wages (QCEW) and the Burning Glass/Lightcast job vacancy data (BG).² We focus on 18 U.S. states with parsable firm-level data on program participation. To understand the need for training subsidies in the private sector, we analyze the characteristics of employer participants and the markets they hire in, relative to employers and markets that have not had grants. To understand how grant receipt impacts labor demand, we examine the impact of program participation on

¹See for instance a McKinsey report (Laboissiere and Mourshed, 2017) which found that “almost 40 percent of American employers say they cannot find people with the skills they need, even for entry-level jobs,” and Forsythe et al. (2022) on the labor supply shortage during the COVID recovery.

²The BG database comes from the company now known as Lightcast. They scrape and code the near universe of job vacancies posted to online websites such as job boards and individual company websites and use proprietary algorithms to parse, deduplicate, and code the content of the ads. See Hershbein and Kahn (2018) for an early use of BG and more details.

employment and vacancies using an event study and nearest-neighbor matching design. Together, these analyses shed light on frictions in the provision of worker skills and assess whether training grants are successful at overcoming these frictions.

These programs typically have explicit goals of helping to upskill the state’s workforce, especially in skills that would be transferable across employers. However, training subsidies may also serve as place-based incentive policies. In practice, we find that grants are much more likely to be used in competitive labor markets, as measured by the concentration of hiring firms or market tightness. Theory suggests these competitive labor markets should experience more underinvestment in training because firms will be concerned about losing trained workers to competitors (workers can face constraints to skill investment everywhere). We also find that grants are disproportionately allocated to larger and higher paying firms and labor markets and to firms seeking to hire more skilled workers. We find no evidence that the grants are used to even out prospects across neighboring markets or that grants are targeted to emerging labor markets, to firms that are new to the state but not new overall, or to megafirms that might have outsized influence across states.

We do not view all firm applications, only the realized grants, so these outcomes are the combined result of firms’ application and states’ allocation decisions. Bureaucratic and monetary costs for program participation (e.g., a requirement to match state funds and raise participants’ wages), make the net benefit to firms unclear. Based on a closer examination of states’ allocation policies and the consistency of results across different regimes, we believe firm choices of whether to apply for these grants drive most of the allocation patterns.

Next, we analyze the impact of program participation on firm recruiting and employment outcomes. We use an event-study design to compare treated firms to a matched sample of similar firms. We use a control group that is matched on several lags of size and propensity to post vacancies captured in Burning Glass, allowing us to compare treated firms to those at a similar observed point in their life-cycle growth and recruiting trajectory. Our nearest-neighbor matching design for firms is in-line with the modern literature on estimating causal effects for workers participating in training programs (Dehejia and Wahba, 2002; Imbens and Xu, 2024).

After grant receipt, vacancies and employment at participating firms increase relative to the control group and continue to grow for many years. Using job vacancies as a proxy for the composition of employment growth, we find that ad shares shift away from professional occupations and toward lower-skilled positions. Perhaps because of these shifts, we see no relative impact on participating firms’ overall wage bill per worker. These effects accrue over time and extend well beyond the training period. This overall employment growth and shift in demand towards lower skilled positions may reflect a complementarity between targeted and untargeted skill groups. Training that

increases the productivity of one group should increase demand for the other as long as they are production complements (as in Katz and Murphy (1992)), while alleviating training frictions can increase employment for the targeted group as well (Restrepo, 2015). Downskilling may also stem from the growth itself as in Engbom et al. (2023), who find that growing firms expand most at the bottom. Finally, these firms may have invested in “training capital” such that they are now willing to take a chance on less skilled workers.

Finally, we explore heterogeneity by the types of skills being trained for, which are available for a subset of our data. We map text descriptions of training programs into broad occupation categories using a large language model (LLM). Our findings suggest two typical archetypes for training firms. The majority of firms with parsable training plans target professional skills, such as leadership and process management. These firms had high demand for professional roles and high skill requirements in their job postings prior to receiving the grant. After training, these firms grow rapidly while reducing skill requirements—as in the aggregate results—and raise wages. The second most common target of these training proposals is production-related skills, with a particular focus on efficient manufacturing and computing technology. These firms had more modest skill requirements prior to training and were not more likely than control firms to post ads in production occupations. After receiving the grant these firms also grow more rapidly, but—unlike the first group—become more likely to hire in the targeted occupation and increase skill requirements in postings. The patterns are consistent with grants supporting growth in two ways. Some firms invest in leadership skills that allow them to grow and hire in complementary occupations. Others use training to pivot their existing workforce, perhaps to work alongside automation technology, creating greater demand for the same occupation.

The evidence we present suggests these grants resolve a skills gap which previously prevented firms from operating at optimal scale. The fact that labor inputs change post grant receipt means that these grants are not purely crowding out private sector funds. Instead, they are, on average, targeting firms on the margin of whether or not to train and facilitating upskilling of the state’s workforce. Grants concentrate in more competitive labor markets, where firms should be most reluctant to pay to train their own workers due to poaching risk. Overall, our findings indicate that Beckerian frictions create genuine underinvestment in worker training. When public funds can overcome these frictions, firms enter a period of prolonged higher growth while expanding opportunities for less skilled workers.

We are the first to provide a broad-based evaluation of these public-private incumbent worker training programs, thereby contributing to a large body of literature on training programs more broadly. A rich literature in economics focuses on government training programs targeted at the

long-term unemployed, or other disadvantaged workers, and tends to be quite pessimistic.³ Card et al. (2018) perform a meta-analysis of a large number of active labor market programs throughout the world and confirm the lack of impact of public sector programs on reemployment, but find positive long-run impacts for other types of programs, such as those in the private sector.⁴ Katz et al. (2022) evaluate a series of sectoral training programs that target training to transferable skills that are in high local demand and find positive earnings impacts. Researchers have highlighted public-private training programs as a potential solution to some of the classic problems with public sector programs, including earlier single-state analyses specifically focused on incumbent worker training programs in Massachusetts (Hollenbeck, 2008), Michigan (Holzer et al., 1993), and New Jersey (Van Horn and Fichtner, 2003). Our systematic cross-state analyses of grant allocation and their impacts help shed light on the motives and benefits of these programs at a broader scale. We are also the first to study these programs following a significant expansion in federal support from the 2014 Workforce Innovation and Opportunity Act.

Our analyses allow us to target the firm as the focal unit of observation. Much of the past research about training at the firm level focuses on how firm-financed training impacts wages and productivity (e.g., Lynch and Black, 1998; Almeida and Carneiro, 2009; Jones et al., 2012; Konings and Vanormelingen, 2015) rather than how these outcomes vary with government subsidies for training. These studies typically rely on survey-based measures of training which are subject to measurement error and yield varying rates of training provision depending on whether respondents are firms or employees (see Black et al., 2023 for a survey of this literature). Our new collection of firm-level data on state provision of training subsidies means we do not need to rely on self-reported training provision, but rather categorize a firm as offering training based on grant receipt.

Our paper also contributes to the literature exploring the relationship between firm-financed training and labor market concentration. Theoretical models (Becker, 1964; Katz and Ziderman, 1989; Acemoglu and Pischke, 1998; Stevens, 1996) predict that there will be under-provision of worker training in more competitive markets due to concerns about poaching. Our paper provides new evidence in the U.S. market that training in the presence of subsidies varies with market concentration. We leverage a growing literature on labor market concentration (Yeh et al., 2022; Berger et al., 2022) and especially those that use BG to measure labor market concentration at a highly disaggregated level (Azar et al., 2020; Schubert et al., 2022). We provide evidence that markets with greater poaching risk may indeed suffer from an under-provision of human capital, thereby

³See for example Ashenfelter and Card (1985); Ashenfelter (1978); Heckman et al. (1998); LaLonde (1986) among many others. These papers tend to find no impacts on program participants and hypothesize that the programs may stigmatize participants, have other close substitutes, face compliance issues, or be poorly run.

⁴O’Connell et al. (2019) compares different types of training programs in Brazil and finds double the reemployment effect for one public employer-informed program compared to a more traditional one.

contributing to a seminal and largely theoretical literature in labor economics on investments in training.⁵

While college graduates learn general analytical skills that help them shift tasks with changing skill demands, a large group of workers with less formal education may instead invest in specific technical skills that can become obsolete. Less educated workers face a risky and unpromising labor market, discouraging them from re-investing in new skills on their own. This uncertainty and rapid change may have opened gaps between the characteristics of the American workforce and the skills employers need now. Our paper provides a better understanding of one policy lever aimed at closing this gap. In turn, our results shed light on the constraints that prevent firms from providing training without public support.

This paper proceeds as follows. Section 2 provides institutional detail on the training programs we study and lays out empirical tests to better understand the use of public funds for private training. Section 3 describes data sources and summarizes characteristics of training firms. Section 4 relates training grants to market-level characteristics. Section 5 examines changes in employment and vacancies as a function of grant reciprocity. Section 6 discusses differences in outcomes based on type of training. Section 7 concludes.

2 Public-Private Incumbent Worker Training Programs

2.1 Policy Context

Public funding for job training programs has existed at the federal level for well over fifty years. However, the majority of this funding – and the majority of researchers’ evaluations of these programs – have focused on funds that target non-employed individuals in disadvantaged groups.⁶ These more traditional job training programs impart skills to the participants that are believed to be valuable in the private sector but typically do not have direct employer involvement. The programs we focus on, in contrast, direct public-sector funds to employers who have applied for a training grant. At the national level, the Workforce Investment Act of 1998 allowed a small use of federal funds for such state-sponsored programs and this allocation was expanded in the Workforce

⁵A number of papers have explored how poaching risk correlates with training provision in the European market, finding mixed support (Muehleemann and Wolter, 2011; Rzepka and Tamm, 2016; Stockinger and Zwick, 2017; Mohrenweiser et al., 2019; Brunello and De Paola, 2008). Our paper provides novel evidence on this question by focusing on the U.S. and specifically tackling the extent to which public-sector involvement can help resolve this friction.

⁶See for instance McCall et al. (2016); Holzer (2023); Hollenbeck and Huang (2015); Heinrich et al. (2013); Andersson et al. (2022) for work on publicly funded workforce development programs more broadly and also Negoita and Goger (2024) for concurrent work on California’s Employment Training Panel with consistent findings to ours.

Innovation and Opportunity Act of 2014 (WIOA). WIOA allows states to spend up to 20% of their allocated federal funds on incumbent worker training grants. These programs have largely been overlooked by researchers since the WIOA expansion.

We conducted a comprehensive search of incumbent worker training programs by browsing state training websites and combing program annual reports for detail. We tracked programs where the primary training grant recipient is an individual firm – rather than a worker or business consortium – to distinguish from traditional training programs. We find that almost every U.S. state has at least one program. States spend, on average, \$0.50 to \$9 per capita on an annual basis. Throughout all analyses in this paper, we will restrict our attention to 18 U.S. states that have parsable firm-level data on program participation (see the map in Figure A.1). We describe the data we collect on these programs in more detail in section 3.

In Appendix section A, we provide a comprehensive comparison of programs, which vary a great deal across states. Here we review some common themes and the most relevant details. In all states, firms initiate the grant application process. Firms must submit a proposal that specifies training needs, a description of the planned training, estimated costs/desired funding, and the number of incumbent or newly hired workers to be trained.⁷ Length of training varies, ranging from under six months to two or three years. Firms can and do apply for new grants once their current grant period is completed; 20% of the firms in our sample have multiple grants.

States typically mention wanting to upskill their workforce and help firms keep up with out-of-state competition. As such, most programs require the training to provide industry-recognized credentials, as well as wage increases for the trainees and a guaranteed retention period. Funding amounts vary with a median of \$1,000 per trainee and a mean of \$2,240. Firms can therefore expect to recoup roughly \$20-\$40 per worker-week but not much of their salary outlay. Instead, money can cover training materials, training infrastructure, and small contributions for the opportunity cost of time. In most instances, firms must provide some amount of matching funds (typically 50% of training costs).

States vary in their allocation processes, with rigorous scoring rubrics and competitive processes in some cases and first-come, first-serve allocations in others. In these latter cases, states may express their priorities through regulatory barriers that influence which firms opt in. Between the limited dollar values, credentialing and pay raise requirements, and administrative overhead surrounding

⁷We include programs that focus on either incumbent workers or on newly hired workers, meaning that the firm can be asking for money with the intention of hiring unskilled workers that will go through the training before starting their job. Conceptually, we consider grants earmarked for incumbent versus newly-hired workers as equivalent. Neither type of grant includes any help to firms in finding workers to employ or any restrictions on who the firm can hire (in contrast with other programs that incentivize hiring the currently unemployed).

these grant programs, we expect substantial self-selection. Firms will likely only apply when they can usefully train a large group of workers and meet the administrative hurdles of application and compliance. In several states, nearly all grant applications are funded, meaning most firms choose not to participate.

2.2 Conceptual Motivations for Empirical Analysis

Firms and states may want to participate in these programs for many reasons. Theoretic motivations for state governments broadly fit into two classes: easing frictions in private human capital provision or place-based development goals. Firms will participate in these programs if they see gains to training their workers and the benefit of the grant money outweighs the administrative and compliance costs of the program. We discuss the state’s motivations next, followed by the firm’s problem, and finally, a roadmap for our empirical analysis.

The canonical theories of human capital investment suggest that employers and employees who have already reached a work agreement should also be able to contract privately to share both the costs for workers to accumulate new skills and the benefits of their resulting increased productivity (Becker, 1962; Mincer et al., 1974). There is no room for the public sector to productively subsidize incumbent worker training. A worker should have to pay the full cost of her training in general skills in a competitive labor market, while the cost of specific skills that are only valuable at the current firm should be split. However, in practice, workers may be reluctant to make these investments due to barriers created by credit constraints (Becker (1964), Belley and Lochner (2007)) and risk aversion (Altonji (1993), Patnaik et al. (2022)). Workers may also face information frictions about the skills demanded by firms or underestimate the potential returns to investing in education and training (see Caliendo et al. (2022) for a review). Small and young firms may also behave like individual workers as they face some of the same borrowing constraints (Banerjee and Duflo (2004), Kerr and Nanda (2009)).⁸

Acemoglu and Pischke (1998) highlight one market imperfection that may solve the underinvestment problem, even in the face of these other constraints. When labor markets are imperfectly competitive, firms can expect to retain their workers and exercise monopsony power. Several recent papers document the degree of monopsony power in many U.S. labor markets (Yeh et al., 2022; Berger et al., 2022). Under monopsony, Acemoglu and Pischke (1998) argue that workers will be

⁸Minimum wage laws can create a further barrier by preventing wages from falling far enough to make training workers in general skills worthwhile for the firms (Hashimoto (1982) and others summarized there), even if workers were willing to incur the cost of training. A large literature has explored the relationship between minimum wages and worker training in practice (see Hara (2017) for a recent survey) with mixed results.

less willing to cover the cost of any kind of training, since their lack of bargaining power will prevent them from extracting the gains of their growing productivity. On the other hand, firms should be more willing to cover the cost of investments – even in general skills – since they can expect to retain the benefits without the threat of poaching. We would expect that if these grants are mainly overcoming under-investment in worker skills due to this poaching externality, then they should be more prevalent in competitive labor markets.⁹

Economic literature also motivates the broader place-based development goals of the state. There is ample evidence that states use incentive programs to compete to bring businesses to their state (Bartik, 2017). Funds earmarked for worker training may be a particularly politically appealing tool to induce a large firm to move or remain in state. These incentive programs may make economic sense for individual states, though recent work estimates only small returns (Slattery, 2025).

If place-based development goals are an important driver of funding, then grants should be allocated wherever the government would like to see growth or employment retention. These may be in areas that are far from the technological frontier where the state would otherwise struggle to attract firms (Neumark and Simpson, 2015), for instance areas with large and healthy neighboring labor markets. As another possibility, a state might offer grants to attract firms to move into the state, in which case we would see grants allocated to establishments that are new in the state, but part of older and larger national firms. Finally, development goals may be targeted towards retaining top employers in the state. Grants would then be allocated to industry leaders or firms with high market shares that are better able to direct funds.

Firms will apply for such programs weighing costs and benefits of participating. As described in the previous subsection, firms face significant bureaucratic costs, such as lengthy applications and documentation of completed training, and monetary requirements, such as the matching of state funds and wage increases for trainees. Firms will therefore likely only participate if they are planning to use the funds to provide training. Participating firms may be on the margin and train only if they receive the grant money, either because their private benefit of training does not exceed the costs or because they are credit constrained and cannot make the optimal training investments without support. If neither of these conditions holds, the grant amounts to a lump sum transfer that subsidizes training the firms would have invested in already. Thus, a key indication of what

⁹Past work also highlights that unions in an imperfectly competitive market can incentivize firms to train workers Booth et al. (2003); Dustmann and Schönberg (2009). Because unions compress the wage structure, training may increase productivity faster than it increases wages, even for general skills, generating positive profits to training. This may explain why we see training grants concentrated in traditionally blue collar industries which have higher rates of unionization than other industries. Other contracting arrangements such as non-compete agreements (Starr et al., 2021) can in theory help firms and workers come to a training agreement but contracting frictions may prevent these arrangements from working in practice.

types of firms accept the grant is whether grants have causal impacts on production.

Effective training should increase the marginal productivity of targeted workers at the firm. This productivity shift could increase or decrease employment of these workers depending on the elasticity of demand for their output. However, the firms' willingness to invest in this training suggests a strong demand for those skills and perhaps unmet needs that the grants can resolve. Restrepo (2015) highlights that with underinvestment in training due to market frictions, unemployment is higher; workers face lower job finding rates due to skills mismatch and firms post fewer vacancies due to the lower expected value of the potential matches. Alleviating these frictions can thus increase employment in part via increased vacancies for training-targeted positions.¹⁰ Finally, if there are complementarities across types of workers, this increase in productivity for the trained workers should increase demand for other types of workers as well, as in Katz and Murphy's (1992) Q-complements.

Therefore, our analysis proceeds in two steps. First, we describe the distribution of grant participants in terms of firm characteristics and labor market features. The patterns of grant allocation can help distinguish between grants that seem likely to be motivated by state-directed place-based incentives, grants to firms who may be motivated by credit constraints, and grants to firms who may under-invest in worker training because they face strong competitive pressures. Second, we will evaluate whether grant recipients in fact change labor inputs following program participation, relative to a matched observational control group. Together, these analyses illustrate whether and how these government dollars help overcome market frictions by funding training when the social benefit outweighs the private cost.

3 Data

3.1 Hand-collected program data

After combing state websites to identify programs that match our criteria, we identified 18 states that not only administer an incumbent worker training program, but also retain and publish data on the specific firms that received grants in at least one year. States vary in the number of years of data available, as well as the information about the training provided. We collect data through 2019 if available and we begin as early as 2002 in California. Appendix figure E.1 provides further details on the availability of grant data by year. In addition to firm name, the majority of in-

¹⁰Though this model does not explicitly include a poaching externality, Restrepo notes that the mechanisms which reduce hiring are exacerbated in environments where contracting on training isn't possible.

sample states also report the county of participating firm, number of trainees requested, and the grant amount. Appendix figure E.2 reports the number and size of grants by state.

For a subset of the states in our sample (California, Kentucky, Massachusetts, New Hampshire, and New Jersey), we have text descriptions of firms’ training plans taken from the grant applications. These descriptions range in length and detail; appendix figure E.3 provides an example of a particularly comprehensive training plan from a company in California. This company manufactures electronic signs and proposes training in machinery as well as a range of basic office skills. We analyze the data on training descriptions separately in section 6.

3.2 Supplemental datasets

We augment our hand-collected information on training grant receipt with data on firm behaviors and outcomes from two sources. The Quarterly Census of Employment and Wages (QCEW) provides administrative data on firm age, industry, employment, and total wage bill. Burning Glass job vacancy data (BG) provides a detailed picture of job posting behavior.

The QCEW is a federal government registry of virtually all businesses in the United States that pay into state Unemployment Insurance programs, plus federal government entities. It covers more than 95% of all jobs and serves as the sampling frame for all Bureau of Labor Statistics establishment surveys. We treat the QCEW as our authoritative benchmark for key firm characteristics and also for determining whether each firm survives from one year to the next.

The BG database of job vacancies is collected by Lightcast, a labor market analytics firm that scrapes websites where job vacancies are posted. Through proprietary machine-learning algorithms, they clean, code, and de-duplicate the scraped ads. Their ad-level data can include the employer name, job location, and job title – which is used by Lightcast to impute an occupation. By targeting over 40,000 websites, the BG data include the near-universe of job openings that are posted online. Their primary business model is to provide analytical tools that help businesses and educators track movements in skill demand. As such, they pay careful attention to measuring the skill requirements specified in job ads. In addition to standard skill measures such as education and experience requirements, they also regularize tens of thousands of key word skills standardized from the open text of job ads. Deming and Kahn (2018) distill these words into a categorization of 10 general skills and show wide variation across firms and geographic space, even within narrowly defined occupations. The data are available consistently from 2010 onwards.

Online job postings are not perfectly representative of all hiring behavior. Previous researchers have found the data to be stable and well aligned with national vacancy trends. Dalton et al.

(2025) match BG vacancies to the QCEW and the Job Openings and Labor Turnover Survey and show how the composition of firms vary across datasets, finding a good deal of alignment, though small and low-paying firms are under-represented in BG.¹¹

We merge grants to establishments in QCEW using firm name, state, and county where available, using a fuzzy match when firms do not have a unique, exact match. We are able to match 95% of grants to an establishment in QCEW. From there, we leverage the QCEW-BG merge from Dalton et al. (2025). 85% of grants in the QCEW sample also have job posting activity in BG in at least one year. The resulting dataset uses firm name-county pairs as its unit of observation – the most detailed level at which we can match. When a firm has multiple establishments in the same county, we consider all establishments to be treated. Throughout, we refer to these name-county pairs as establishments or firms, although the precise unit of analysis is sometimes somewhere between the two. Further detail on the matching process is described in Appendix Section B.

3.3 Characteristics of Training Firms

We use the QCEW and BG samples to form a comparison group of firms. To begin with, we restrict attention to the universe of establishments in states and years in which grant data are available. For grant firms, we restrict our sample to firms receiving grants post-2010 and use the first observed grant as the focal year of grant receipt. For each non-grant firm, we randomly assign a “placebo” grant year to match the empirical distribution of actual grant years in the state. From here, we restrict attention to grant and non-grant firms that have non-zero employment in the year of grant receipt (or placebo year) and in the prior year.¹² This placebo year assignment will help us select a time window to compare treatment and control firms and for sample selection criteria in our analyses.


The resulting sample includes 8,667 grant firms and 1.7 million non-grant receiving firms. Appendix table E.1 provides summary statistics for grant firms (column 1) and this full set of control firms (column 2) measured in the year before grant receipt (or the placebo year). We provide summary statistics for both the matched QCEW sample and the set of firms that ever post in BG. We also summarize differences across groups in the distribution of firm characteristics in Figures 1 and

¹¹See also Hershbein and Kahn (2018) who use the BG micro data to understand how the Great Recession changed demand for worker skills. They include a wide range of sanity checks on the data and BG has since risen in popularity among academics.

¹²Most of the time, the restriction on non-zero employment helps us focus on firms that are in operation during the grant time window. However, due to data noise issues, some firms are observed with zero employment for random years in the middle of their spell of operation. In analyses below, we drop these years. Also, from the initial set of firms, we exclude those with no more than 1 employee for average monthly employment, as this group of firms is highly unusual within grant-receiving firms but represents a non-trivial fraction of establishments.

2. These figures take the share of grant-receiving establishments with a given characteristic (for instance, size bin or industry) and subtract the non-grant recipient group share.

Figure 1: Treatment-Control Differential in Distribution of Establishment Characteristics



Notes: We plot the difference between the fraction of treatment establishments in a bin and the fraction of control establishments. We do this for characteristics in the year prior to grant receipt (or placebo year). Wages are quintiles of total wage bill per worker. Growth rate is measured as the t-2 to t-1 change in employment. Any BG Posts tracks the share of firms with ads captured in Burning Glass in t-1. The ad share across occupations in BG restricts to firms that post ads in t-1. See footnote 13 for definitions of the broad occupation categories.

Beginning with the employment size distribution in the top panel of figure 1, we can see that grant recipients are substantially less likely – nearly 50 ppts – to be in the smallest size class (less than 50 employees) and substantially more likely to be among the middle size classes (especially 50-249). On average, recipients had 225 workers, compared to the control average of about 20 (appendix table E.1). Training is likely cheaper per worker at large firms that benefit from returns to scale (Barron et al., 1987). The smallest firms may be unable to make training worthwhile, even with partial subsidies. Interestingly, while grant firms are larger on average, we do not see an

overrepresentation among megafirms (5000+ employees). Grant recipient establishments are older (by about 3 years on average), with substantial overrepresentation (17 ppts) among the oldest bin (11+ years). There are fewer treated firms that were brand new upon grant receipt than in the control group.

We only observe total wage bill in the data, so we define wages here as the total wage bill in a given quarter divided by the number of employees on the first day of the quarter. For many reasons, this payroll per worker metric is not equivalent to average wages. As such, we report a coarse grouping in figure 1, splitting firms into quintiles. Grant recipients are substantially less likely to be found in the lowest wage bin (37 ppts) and much more likely to be found in the middle and high wage categories. On average, grant firms' payroll per worker is about \$20K more than the control group (appendix table E.1).

Grant recipients are less likely than control firms to be shrinking by more than 10 percent of their employment and more likely to be growing at a moderate rate (i.e., 2 to 10 percent). Grant firms are also much more likely to be recruiting online prior to grant receipt – 51% can be matched to BG in year t-1, compared to only 16% in the control group. Consistent with their faster growth, grant recipients post substantially more ads than the control group, even conditional on postings any ads – from appendix table E.1 they average over five times more ads posted in the year prior to grant receipt. Panel B of the same table also shows that, within the BG sample, differences in establishment characteristics across grant and non-grant recipients are similar to those in the full sample.

The ad characteristics provide a sense of the skill level of desired workers for grant versus control firms. First, BG codes whether employers specify an education requirement or a requirement for experience in the field, and, if so, how many years. As shown in Panel C of table E.1, grant firms specify skill requirements at higher rates, including education (experience) requirement in 68% (55%) of ads, compared to 54% (48%) in the control group. Treated firms are also more likely to require a college degree (a subset of all education requirements).


Consistent with their higher skill requirements, treated firms hire in more skilled occupations. We use a coarse grouping of four broad occupation categories: Professional occupations are high-skilled white collar positions; administrative occupations are routine white collar positions (such as sales and office support); service occupations are low-skilled positions like servers and personal care jobs; production occupations are blue collar jobs.¹³ We also use these categories when measuring firms'

¹³This grouping maps SOC occupation codes into four mutually exclusive and exhaustive SOC occupation code groups: Professional includes SOC 11-19, 23, 27, 29; Administrative/Sales is 21, 25, 31, 41 (excluding 412), 43; Low-skill Service is 35-39, 412; and Blue Collar is the remainder (33, 45-53).

recruiting behavior. Among the firms who use online hiring services, treated firms have a greater proportion of job ads asking for professional skills (67% relative to 47%) and are less likely to be searching for skills relevant to administrative/sales, service, or production occupations.

Finally, grant recipients are concentrated in different industries than non-recipients. From figure 2, we find grant recipients are more likely to be in manufacturing industries, whereas non-grant recipients are more likely to be in services (such as accommodation and food, professional and business services, and retail trade). Though not shown, the differences across grant and non-grant groups shown in Figure 1 persist even after controlling for the differences in industries shown in Figure 2.

Figure 2: Industry Distribution across Grant and Non-Grant Recipients



Notes: We plot the difference across treatment and control group in the share of establishments in two-digit NAICS sectors.

4 Grant reciprocity and labor market characteristics

Grant distribution across labor markets reflects the joint outcome of firm applications and state allocation decisions. States cite many different priorities for their programs including a desire to reach small firms, under-served locations, and firms struggling to keep up with out-of-state competition. States also express a desire to provide workers with industry-recognized skills that employers may not be able to find or fund on their own. Economic theory suggests that more competitive labor markets, where poaching risk is greatest, will have a greater need for this type of government intervention. However, as discussed in section 2, not all states allocate grants through competitive or strategic processes. In these cases, the distribution of grant recipients will be driven primarily by which firms choose to apply, which may or may not align with the firms that represent the greatest social return to training grants. For instance, we have already seen that grants disproportionately serve larger, faster-growing, and older firms, despite the fact that several states express a preference for small businesses. Given the high administrative barriers, the relative strength of public priorities and firm needs in determining the distribution of grants is therefore an empirical question, which we tackle next.

4.1 Methods

In equation 1, we relate the likelihood that a labor market receives a grant in a given year, t , to a vector of market-level measures of economic activity motivated by our discussion above. The regressors are defined using a benchmark time period (2010-12) and we explore grant allocations in subsequent years (2013-19).¹⁴ We limit this allocation analysis to state-years in which we observe grants. Markets are defined by commuting zone, c , and skill, j , which we classify by either occupation or industry. Our baseline specification controls for commuting-zone-by-year fixed effects ($\theta_{c,t}$) to adjust for geography-specific time trends across all industries to adjust for the relationship between economic activity and state’s grant allocation within the specific grant cycle.

$$Grant_{cjt} = \beta_0 + f(concentration_{cj})\beta_1 + \mathbf{X}_{cj}\beta_2 + \beta_3 NewMarket_{cj} + \theta_{c,t} + \theta_{s,t} + \theta_j + \varepsilon_{cjt} \quad (1)$$

We also control for state-by-year fixed effects ($\theta_{s,t}$) because some commuting zones cross state borders and for skill (θ_j) fixed effects to adjust for industry- or occupation-specific grant propensity. We cluster standard errors by state to account for persistent state-level correlations in grant allocation decisions.

¹⁴We choose these years because 2010 is the earliest year for which we have consecutive coverage of the Burning Glass data, which we will use to measure market concentration.

We add measures of economic activity that align with the motivations discussed in section 2. To understand poaching risk, we follow the previous literature in defining measures of market-level concentration of vacancy postings using Burning Glass (Azar et al., 2020). Our preferred measure of labor market concentration is a Herfindahl–Hirschman Index (HHI) for job vacancies as in equation 2, calculated using the full universe of job ads posted in BG from 2010 to 2012.

$$HHI_{cj} = \sum_k \left(\frac{(\# \text{ of ads})_{kcj}}{(\# \text{ of ads})_{cj}} \right)^2 \quad (2)$$

The HHI in market cj is the sum of squared ad shares across all firms, k , posting in the market. A higher value indicates that a greater proportion of job vacancies in a given market are from a smaller number of firms (i.e., a less competitive market). This vacancy-based market concentration measure is particularly salient for thinking about poaching risk.¹⁵ Our preferred labor market measure is defined at the two-digit-industry-commuting zone level. We show that our results are robust to defining markets by 3-digit occupation and to alternative measures of poaching risk.

We also explore the relationship between grant receipt and a range of other market characteristics ($\mathbf{X}_{cj}$) such as size and average wage, based on American Community Survey (ACS) data, as well as growth in these measures.¹⁶ We also use the ACS to measure CZ-wide unemployment rates. To better understand economic activity in neighboring markets, we also calculate “leave-out” versions of these measures at the state-industry or state-occupation level (omitting the focal CZ-skill market from that calculation) and the Census division-skill (omitting the focal state from that calculation). We are therefore primarily capturing the relationship between grant allocation and persistent, historical economic health, rather than year-to-year fluctuations.

We restrict measures of economic activity to CZ-skill pairings which have at least 50 ads posted in the base period (2010-2012), ensuring these markets have enough active employers to reasonably measure recruiting. However, we would like to explore whether grants are allocated to markets with little past activity, consistent with a place-based incentive policy designed to draw in large firms from out of state. We therefore include these markets in the regression with the indicator $NewMarket_{cj}$ and impute values of zero for all measures of baseline economic activity.


¹⁵BG measures of market concentration are valid for our purposes since Dalton et al. (2025) show a strong cross-sectional correlation between these and measures that correct for mis-aligned firm composition in BG, even while BG overstates the level of vacancy concentration due to an underrepresentation of small firms.

¹⁶We use ACS 2010-2012 waves (Ruggles et al., 2022), combined with crosswalks between public-use micro areas from Dorn (2009). We calculate the average number of employed people age 25 to 64 working in each market per year and the average wage per hour for workers in this age range in each market. For growth, we use the change in these variables between 2010 and 2012. Average wage is defined as the total earnings from wages and salary, divided by the reported usual hours worked per week times weeks worked in the past year. We top- and bottom- code wages, omitting individuals whose reported salary and hours worker indicate an hourly wages less than 5 or more than 150 dollars per hour.

4.2 Results

Figure 3 provides a bin scatter of the likelihood that an industry-CZ market receives at least one grant on the y-axis and the market-level HHI on the x-axis. We see that lower HHI (i.e., more competitive) markets are more likely to receive grants. The relationship is non-linear, quite steep in the beginning and flattening for higher levels of concentration. This pattern motivates the quadratic functional form we will use in our regression analyses.

Figure 3: Training Grants and Market Concentration



Notes: We divide markets (CZ-by-two-digit industry pairs) into 20 equally-sized bins based on the HHI of job vacancies posted from 2010-12 (see equation 2). For each bin, we then plot the average HHI and the share of markets that received any grants between 2013 and 2019.

Appendix figure E.4 shows a similar relationship with concentration for the total number of grants or grant dollars (including zeros) allocated to a market, so for remaining analyses we proceed with the indicator for whether the market ever received a grant. This simple negative relationship is suggestive of the theoretical mechanism described above where markets with greater poaching risk face an underprovision of general skills. However, less concentrated markets may receive more grants for reasons other than market concentration. For instance, larger markets may be less concentrated and would also mechanically receive more grants even if grants were randomly allocated across firms.

Our multivariate analysis, reported in Table 1, controls for size and many other possible drivers of

grant allocation. Column 1 shows that the negative relationship between HHI and grant receipt holds after controlling for pre-period, market-level employment and wages, CZ-year unemployment, and two sets of fixed effects. Industry fixed effects control for the possibility that certain industries are in favor with state grant agencies and also happen to be more or less concentrated. CZ-by-year effects control for trending differences across CZs that might impact grant allocations. Because some commuting zones cross state lines, we also include state-by-year controls to place our comparison within a grant cycle. To provide some context for the magnitude of the relationship, the mean and standard deviation of the HHI are 0.155 and 0.150, respectively, and the average market receives a grant with 20.5% likelihood, meaning that a one standard deviation increase in HHI from the mean is associated with a 3.0 ppt (15%) decrease in the likelihood that a market receives a grant.

We also see evidence that grants are more likely to go to stronger labor markets, in terms of number of workers, average wages, and the unemployment rate. A market with 1,000 more workers is associated with a 2.9 ppt (14%) higher likelihood that the market received at least one grant; CZ-industries with a 1 ppt higher unemployment rate are slightly less likely to receive a grant (by about 0.4 ppt).

Column 2 of Table 1 illustrates robustness to the inclusion of industry-by-year and industry-by-state fixed effects. The former helps if there are any skills that are rising in popularity that happen to have more or less concentrated markets on average, for instance, states may increasingly value programming skills and jobs in the technology industry may tend to be located in concentrated markets. Industry-by-state effects help control for the possibility that preferences for a given industry are clustered in particular states that also tend to have more or less concentrated markets, for instance, California may preference programming skills and Silicon Valley may be an especially dispersed market. Reassuringly, the negative relationship between HHI and grant allocation holds within these controls.

Columns 3 and 4 test whether grant receipt is associated with the economic characteristics of the surrounding region. We control for own-market employment and wage growth, as well as neighboring market employment, wage level, and growth rates. Column 3 defines the neighboring market as the population-weighted average of all other industry-CZ's in the state, the "leave-out state" market, while column 4 uses the population-weighted average of all other state-industries in the census division, the "leave-out region" market. We revert to the original sets of fixed effects since these neighboring market variables have little or no variation within the richer fixed effects. We find little evidence that characteristics of the industry within the state as a whole impact the empirical grant distribution, nor do neighboring states.


Table 1: Training firms and market characteristics: 2-digit Industry-by-CZ

Dependent Variable	Any Grants Received (mean = 0.205)			
	(1)	(2)	(3)	(4)
HHI	-0.457*** (0.101)	-0.388*** (0.082)	-0.444*** (0.104)	-0.463*** (0.105)
HHI ²	0.559*** (0.118)	0.498*** (0.108)	0.544*** (0.116)	0.566*** (0.121)
CZ unemp rate	-0.338** (0.156)	-0.351** (0.159)	-0.330** (0.155)	-0.340** (0.157)
New Market	-0.024 (0.030)	-0.028 (0.024)	-0.028 (0.034)	-0.023 (0.035)
Employment (1,000s)	0.029*** (0.004)	0.025*** (0.003)	0.028*** (0.004)	0.029*** (0.004)
Wage (\$100s)	0.257** (0.119)	0.132 (0.097)	0.176 (0.179)	0.260** (0.113)
Emp growth			0.006 (0.018)	0.013 (0.019)
Wage growth			0.006 (0.050)	0.022 (0.046)
Leave-out State Emp			9.054 (8.546)	
Leave-out State Wage			0.016 (0.269)	
Leave-out State Emp Growth			0.070 (0.072)	
Leave-out State Wage Growth			0.233 (0.169)	
Leave-out Region Emp				-0.224 (1.622)
Leave-out Region Wage				-0.003 (0.024)
Leave-out Region Emp Growth				0.007 (0.107)
Leave-out Region Wage Growth				0.046 (0.122)
State-Yr, CZ-Yr, Ind. FE	X	X	X	X
Ind-Yr, Ind-State		X		
Observations	20,031	20,030	20,031	20,031
R2	0.385	0.428	0.386	0.385

Notes: Standard errors in parentheses clustered on state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Observations are 2-digit industry-by-CZ-by-year. HHI, Employment, and Wages are industry-by-CZ averages from 2010-12. Emp and wage growth are the rate of change in 2012 from 2010 for the industry-by-CZ. The CZ unemployment rate varies by industry and year. The Leave-out State and Region variables are also at the industry-by-geography level, averages over 2010-12 or the rate of change over that period and leave out the focal CZ or state, respectively. Regression observations restricted to 2013-2019. Covariates are defined for the 2,255 markets that posted at least 50 ads in the baseline 2010-12 period and have coverage in the ACS, and other markets are considered “New”.

Finally, we see that “new” markets – those that rarely show up in the vacancy data – are consistently less likely to receive grants in the multivariate analysis, though these are noisily estimated. Figure 4 provides bin scatters of number of ads posted in the baseline period (2010-12) including zeros and the likelihood that the market receives any grants (left) or average grant size in the market (right, including zeros) in the analysis period (2013-2019). These plots find the consistent pattern that markets with more posted ads are more likely to receive grants. In other words, grants are observed across markets in proportion to their baseline economic activity. We do not see evidence indicating other types of strategic choices by states, for instance disproportionately targeting new markets which might be the case if grants were frequently designed to entice large firms to relocate to new markets.

Figure 4: Training Grants and # Ads Posted in the Market



Notes: We divide markets (CZ-by-two-digit industry pairs) into 20 equally-sized bins based on the number of ads posted in the market from 2010-12. We then plot the share of markets that receive at least one grant (left) or the average dollar amount per grant (right, including zeros) on the average number of ads in the baseline period. We restrict to markets that post no more than the 90th percentile (1277 ads) for visual clarity for the smaller markets, though the slope of the line is fairly similar when we include them.

Appendix figure E.5 and table E.2 show the results hold when defining markets by CZ and occupation rather than industry. The goal in defining these markets is to identify a specific skill that an employer might wish its employees to have and better understand the labor market prospects for that skill. Because grants are allocated to firms, not occupations, and the QCEW provides information on industry of the firm but not occupational mix, our primary measure is based on

industry. However, because occupation is a more natural analog to the skills that define a worker’s outside option, we also use the ad distribution of the establishment to allocate grants to the modal occupation among the firm’s job postings. Reinforcing the findings at the industry level, grants are significantly more likely to be allocated to bigger and more competitive markets.

Finally, we explore other measures of market competition, rather than the HHI of job vacancies. Results reported in appendix table E.3 show that the result that grants are more common in more competitive markets holds up when considering an alternative functional form for concentration: the share of ads posted to the three largest firms in the market (defined as either industry- or occupation-location). When markets are defined by industry, we can also explore the concentration of employment shares using County Business Patterns. In both cases, more concentration of ads/employment is associated with lower grant receipt. Lastly, we show robustness to another measure of market competition: labor market tightness. Tighter markets should have more poaching and indeed we find that they are also more likely to receive grants.¹⁷

4.3 Discussion

We find a strong and robust negative relationship between market concentration and grant allocation. This pattern is consistent with firms that are reluctant to pay to train workers when they are competing heavily for talent within their market. In these instances, public sector subsidies can help solve market failures. It is not clear from these results whether the correlation is driven by state governments targeting grants to competitive labor markets or firms in these markets applying at higher rates. To better understand these trade-offs, we categorize states based on whether the grant allocation process seems to be competitive (i.e., a strategic evaluation process that results in only some applicants receiving grants based on public priorities) or firm-led (e.g., first-come, first serve) and look at whether the relationship between HHI and grant receipt varies by allocation method. For the latter, allocations will be driven almost completely by firm application decisions, while, in the competitive case, allocations will be driven by the combination of firm application decisions and state allocations. Appendix Figure E.6 shows a bin scatter of the likelihood of grant receipt against HHI separately by state-level allocation method. We see that both types of states have a similar negative relationship between concentration and grant receipt. This similarity suggests a strong role for firm application decisions in driving the empirical correlation.

¹⁷We define tightness at the CZ-industry level as the average number of vacancies posted in BG between 2010-2012 divided by the average number of unemployed workers who previously worked in the industry as measured in the ACS over the same period. Another desirable measure would be the rate of job-to-job transitions in the labor market, but it is unfortunately not possible to measure transition rates at these levels of granularity.

We also see that grants are allocated to bigger, well-established, higher paying markets, with lower unemployment rates. Grants are not allocated to new markets or markets with growth capacity (i.e., small and fast growing). If anything, grants are instead allocated to markets whose neighbors exhibit growth capacity. These patterns would seem to be at odds with place-based development policies that may prioritize markets that are lagging their neighbors or typically prioritize small or growing markets. In section 3, we also saw that grants are allocated to older, larger, and faster growing firms. The fact that grants go to more established firms and markets could be evidence of regulatory capture, though we do not see that grants are more likely to go to industry leaders or firms with very high market shares themselves. Furthermore, if place-based policies targeted large firms that had greater regulatory capture, we might have expected the allocation to go towards more concentrated markets overall.

Grants are not free money. Firms face both administrative costs when applying for these grants as well as financial costs in the form of matching funds and guaranteed wage increases for trainees. Hence, only some firms will find the program worthwhile to participate in. Our finding on concentration is consistent with the self selection driven by firms whose social value of training is larger than its private value due to poaching risk.

5 Outcomes of Grant Recipients

Having established that grants tend to concentrate in more competitive labor markets, we next turn to the question of whether individual establishments change their employment and hiring behavior in response to receiving a grant.

5.1 Methods

We estimate a series of event study models, leveraging two-way fixed effects to compare the firm-level outcomes for grant recipients to the trajectory for non-recipients. Equation 3 specifies a regression of outcomes for firm i in year t on an indicator for whether t is τ periods before or after the grant year, T , defined as the first year we observe the firm receiving any grant. Because we have assigned placebo training years to the control group, we can also control for placebo event time (i.e., main effects in event time), which can help to address problematic control comparisons that may arise in some specifications with staggered adoption of treatment (Sun and Abraham, 2021; Goodman-Bacon, 2021). The event time indicators of interest are all interacted with the ever treated indicator $-\mathbb{1}(grant_i)$.

$$y_{it} = \beta + \sum_{\tau \neq -1} \beta_{\tau}^{grant} \mathbb{1}(t = T + \tau) * \mathbb{1}(grant_i) + \sum_{\tau \neq -1} \beta_{\tau} \mathbb{1}(t = T + \tau) + \theta_i + \theta_t + \varepsilon_{it} \quad (3)$$

We restrict attention to grants received between 2010 and 2019 and restrict the regression sample to a window surrounding grant receipt (or placebo receipt) of at most plus or minus 5 years. For outcomes measured in the QCEW, we begin the sample as early as 2005 – to observe a full five years pre-treatment for even the earliest treated cohorts – and stop our analysis in 2022 due to data availability. BG data are only available from 2010-2022 so the earliest treated cohorts are not observed in the pre-period. Our baseline sample is therefore an imbalanced panel.¹⁸ We explore a wide range of outcome variables to better understand patterns in labor inputs. These include log employment and wage bill per worker as measured in the QCEW, the number of vacancies posted in BG, and the distribution of vacancies across occupation groups and skill requirements.

As illustrated in the previous section, firms that apply for and receive training grants are different than the typical firm in substantive and important ways. The key two-way fixed effects identifying assumption of parallel trends is therefore unlikely to hold in the full sample. To address these asymmetries, we take lessons from the classic literature estimating the causal effects of training programs for workers, which has repeatedly emphasized the importance of matching or other techniques to ensure overlap in the characteristics of treated and untreated units (Dehejia and Wahba, 2002; Imbens and Xu, 2024).

Firms may choose to apply for a grant at a particular point along their life-cycle where training is especially valuable. With or without the grant, such firms might invest in workers or infrastructure and experience commensurate growth and other changes, creating a threat to internal validity. Our approach assumes that these life-cycle dynamics are reflected in pre-period growth and recruiting behavior. We use a matching design to select a group of control firms that track the pre-grant development of the treated firms. We leverage the scale of our data sample to match directly on multiple pre-treatment (or placebo treatment) growth, as in Abadie et al. (2010) or Arkhangelsky et al. (2021a), rather than collapsing these trends into a single propensity score. Specifically, we identify a single nearest-neighbor (Rosenbaum and Rubin, 1983; Abadie and Imbens, 2006) in the same industry as each treated firm that minimizes the Euclidean distance between the treated and the control firms on 1) five lags of log employment leading up to treatment (or placebo treatment) and 2) five lags of the number of BG postings.¹⁹ To avoid capturing spillover effects within our control sample, we exclude all untreated firms in industry-county markets where at least 20% of

¹⁸Balanced panel estimates, with a much reduced sample size, are qualitatively similar but noisier.

¹⁹For supplementary analyses which use BG data only, we match on only the five lags of number of ads, including zeros.

workers were at firms that received a training grant in any year within two years of the (placebo) treatment.²⁰ Following Abadie and Spiess (2022), we match without replacement, which allows us to construct valid analytic confidence intervals in the later event study regressions by clustering on match pair.

Appendix Section C describes the matching algorithm in more detail. We follow the recommendations of, for example, Crump et al. (2009) and Young et al. (2026) to restrict our sample to treated firms for whom a strong control match exists (dropping 1% of treated firms). We also discuss placebo tests to validate our specification and robustness to alternative matching methods (propensity score matching and synthetic difference-in-differences).

Figure C.1 plots the mean values of the targeted outcomes for the grant-receiving firms and the selected matched control firms. After matching and trimming, both treated and control firms exhibit little to no growth in log employment in the pre-period and strong, parallel growth in number of posts. For both outcomes, the treated firms show a clear trend break after grant receipt while the control firms do not. This figure provides reduced-form evidence that grants do impact the trajectory of firms and that our estimated effects are not driven by mean reversion within the control group.

The alignment of treated and matched control firms is also apparant in Table E.1, which provides a comparison of treated and control firms in the matched sample (columns 3 and 4) to treated and control firms in the full sample (columns 1 and 2). While the full set of non-grant firms is smaller, lower paying, and younger than treated firms, the matched sample is much closer on these dimensions. By design, the matched sample is also quite a bit closer on the propensity to post vacancies in BG, which helps not only conceptually – since we compare firms with similar hiring needs in the pre-period – but also with later analysis on the composition of postings that must restrict to firms that post ads in BG. Wage bill per worker and the distribution of ads in BG are not targeted, but treated and control observations also look similar on these variables in the matched sample.

For regression analyses, we add controls for establishment fixed effects (θ_i) to absorb any time-invariant differences across program and non-program participants and calendar year fixed effects (θ_t) to absorb any common macroeconomic shocks.²¹ We showed in the previous section that differences across grant and non-grant firms are interesting in their own right to better understand

²⁰Excluding all markets that ever receive a grant would eliminate nearly all potential control firms in large, urban markets. This 20% 2-year cutoff strikes a balance of eliminating markets with the strongest potential spillovers while allowing us to construct a sample of well-matched comparison firms.

²¹In robustness checks, reported in Appendix Section E, we add controls for industry-year to capture common sectoral shocks. Results are not substantively different though more noisily estimated.

the nature of the skills gap problem. This section holds constant these differences and allows us to explore differential labor impacts post training. The matched control approach provides the cleanest estimate of the added impact of these training grants apart from selection effects.

5.2 Results

Quantity of Employment and Vacancies


We begin by considering log employment as measured in the QCEW. Figure 5 plots event study coefficients and 90% confidence intervals. Appendix Table E.4 reports the coefficients and standard errors for this specification, as well one that adds sector-year fixed effects. Firms grow rapidly after receiving their first training grant. By five years after grant receipt, treated firms have grown roughly 12.4%, relative to their counterfactual trajectory.

The mean grant receiving firm in the match sample had 186 workers in the year before grant receipt (table E.1). For this mean firm, 12.4% growth represents an additional 23 workers over five years. When firms specify a number of new hires to be trained (40% of the time), that averages to 58 additional workers. However, those new hires would come at the beginning of the grant period, while the bulk of excess growth in Figure 5 accrues only over time. Grant amounts average roughly \$90,000, though factoring in that many firms receive multiple grants, the total grant dollars received by the average firm over 5 years is \$150,000. Furthermore, most firms are required to match government fund 50-50, for a rough total investment of \$300,000. Still, the growth we document here, especially outside the initial grant period, is large. Were the 23 additional workers drawn from non-employment, this policy would be an extremely effective use of government funds for generating job growth. However, we have no way of knowing whether these workers were drawn from non-employment or from other jobs.²² We are unaware of any other study that measures the impact of these types of programs on measured employment, but these large changes are broadly consistent with substantial productivity changes measured by the small literature on the impacts of firm-led worker training (for example, Konings and Vanormelingen (2015) estimate that firm-led training in Belgium makes workers 23% more productive on average).

We next examine the quantity of vacancy postings to better understand whether firms had a demonstrated preference for this growth (as opposed to passive hiring or reduced attrition). The second panel of Figure 5 and the third and fourth columns of Appendix Table E.4 report effects of

²²We conducted some exploratory analyses at the market level and found noisy zero impacts on overall employment and churn following receipt of a grant by at least one firm in a market. The small scale of these grants relative to the size of the typical treated market makes it unlikely that would be able to detect any market-level effects.

Figure 5: Firms Grow After Training: Employment and Vacancies Event Studies



Notes: This figure reports coefficients estimated using equation 3, our event study regression, for log employment measured in QCEW (left) and number of postings in BG (right). We limit to firm-year observations with non-zero employment and the right panel further limits to years of BG availability. Within these conditions, we impute a value of zero postings if the firm did not show up in BG in that year. We control for establishment fixed effects, year fixed effects, and event dummy main effects. We plot coefficients on event dummies interacted with treatment. We also report the 90% confidence interval for standard errors clustered at the match pair level.

receiving a grant on the annual number of BG posts. For this sample, we restrict to years where BG is available and to firm-year observations with non-zero QCEW employment. If a firm meeting this restriction does not post in BG in that year, we impute a value of zero. By leveraging the information in QCEW to condition on a firm being operational, we believe a lack of ads in BG implies zero posted vacancies.

Effects on the number of postings are consistent with the net changes in employment, with steady growth building for many years after grant receipt. By five years out, the average treated firm posts 10 more ads per year than the matched control firms. Almost all grants last for two years or less, with most being completed within a year of receipt. Therefore, while some firms may have increased hiring needs around the time of grant receipt due to a promise to train newly-hired workers, the mechanical effect cannot explain the increases in the later years shown.

Composition of Vacancies

If training grants allow firms to address skill shortages, we expect the composition of whom they hire may change following grant receipt. To explore employment composition, we exploit the rich detail in the BG vacancy data, using job ads as a proxy for how the workforce is changing. For the following analyses, we use a data set which uses only BG information to match grants to firms and to match treated firms to their nearest neighbors.²³ As such, our analysis restricts to firm-year observations with non-zero BG posts.²⁴


We start by examining the occupation distribution of the expansion in postings. To create a single measure for occupational skill, we assign each 2-digit occupation its average hourly wage as measured in the Occupational Employment and Wage Statistics (OES) in 2010. The vast majority of job ads do not post a wage, so assigning average occupational wages is a general proxy for the skill level of the job ad – and not an exploration of whether the firm pays idiosyncratically high or low wages. We aggregate ad-level observations to the firm-year level, the level of variation underlying our covariates. We then regress ad-weighted average occupational wage on grant receipt timing and weight firm-year observations by the number of posts, thus making these equivalent to ad-level regressions. Figure 6 and Appendix Table E.5 show that as firms expand employment and postings they downskill, shifting postings towards occupations with lower typical wages. The change emerges gradually, in line with the other outcomes, and is modest but statistically significant. By four years post grant, the occupations demanded by treated firms are almost \$1.5 less skilled (as measured by occupational hourly earnings differentials). This reflects a 5% decline from the pre-period mean or about 23% of a standard deviation in this measure. Appendix figure E.8 shows that this aggregate decline is driven by a shift away from professional occupations post grant receipt and towards mid-level administrative and production occupations.

These changes in the occupation composition of vacancies occur in the context of strong overall growth in log employment and total job posts, so the decline in the *relative* frequency of professional vacancies reflects a less-than-proportional growth in hiring in this space rather than an absolute decline. As we will discuss below, a plurality of firms with recorded training plans target skills for

²³See appendix B for details on the outside-BLS sample using the same name-matching algorithm as above. This choice is due to constraints which made additional disclosures of internal BLS results impractical. For the BG only NN match, we match on 5 lags in number of postings, which is analogous to our match on firm size, and follow all other matching protocols.

²⁴For comparability, Appendix Figure E.7 reproduces QCEW results for log employment restricting to the sample of firm-years with at least one BG ad. We explored a range of restrictions on the number and regularity of job postings to ensure that the posting distribution across occupations and skills captures a large fraction of firm activity. Because these regressions are weighted to focus on the average posting (as opposed to the average firm), such restrictions make almost no difference.

Figure 6: Average Occupational Wage of Job Postings (BG)




Notes: See figure 5 for regression specification information. Here we restrict to establishments that had BG postings in the year. The outcome is the average occupational wage assigned to 2-digit occupations using OES and weighted within a firm based on the occupation's ad share. Bars indicate 90% confidence intervals with standard errors clustered on matched pair. This figure uses the BG-only sample; see appendix B for details.

professional workers. Therefore, we find that firms who receive training grants grow faster overall and shift modestly away from hiring in the most commonly targeted occupation and towards all other groups.

BG also allows us to explore a wide range of skill requirements, one of the most common being a college degree. The solid line of Figure 7 (and Appendix Table E.5) shows that firms reduce their relative demand for college-educated workers following the receipt of their first grant. By four years after grant receipt, the share of vacancies requiring a college degree or more declines by about 7 ppts (21%). This shift can be entirely explained by changes in the occupations firms are hiring into. As illustrated by the dashed line in figure 7, the change in college degree requirements is indistinguishable from zero after adding controls for ad shares in the 4 broad occupation groups.

A college requirement is one of the most verifiable descriptors in a job ad, and a common barrier to entry for unskilled workers (Fuller et al., 2017), so we were interested in this outcome ex ante. However, we should caution that we found less movement in a range of other skill requirements. Appendix figure E.9 summarizes event studies for any education, any experience, and two of the

Figure 7: College Requirements (BG) Event Studies



Notes: See figure 5 for regression specification information. Outcome is the proportion of ads specifying a college degree or higher. Regressions with dashed lines control for the composition of ads across the 4 broad occupations groups in the year. Bars indicate 90% confidence intervals with standard errors clustered on matched pair. This figure uses the BG-only sample; see appendix B for details.


most common keyword skills categorized in Deming and Kahn (2018). On average, we do not see systematic changes after training. This null effect on any education requirement implies that firms shift requirements from college to high school, which is the other most common listing. That movement could generate a significant change in accessibility (actual or perceived) to a firm. As we discuss in the next section, these null average effects mask substantial but heterogeneous changes across firms that target training to different types of workers.

Wages

The previous section documented the impact of training grants on predicted average wages based only on the occupation mix of job postings at the firm. Realized wages may behave differently at treated firms because training will increase the marginal productivity of targeted and incumbent workers. In fact, several states require that firms achieve certain pay raises for the trained workers as a condition of the grant. Several earlier studies (Jones et al., 2012; Konings and Vanormelingen, 2015) have found that firm-led training increases wages for the targeted workers, though growth in

wages tends to be substantially smaller than growth in worker productivity. We explore changes to firms' total wage bill using the QCEW-match sample, though we have no way to distinguish the wages of workers targeted for training. As shown in Figure 8 and Appendix Table E.4, we find no significant changes in the log total wage bill (controlling for log employment) following receipt of a training grant.

Figure 8: Log Wage Event Studies



Notes: This figure reports coefficients estimated using equation 3, our event study regression, for the log of total wage bill in the year, controlling for log employment. In the BG postings sample (restricted to firm-years with non-zero postings), we run a specification without (hollow circle) and with (hollow triangle) controls for proportion of ads in the 4 occupation groups in the year. Bars indicate 90% confidence intervals with standard errors clustered on matched pair.

Several factors may explain these non-effects. First, and most importantly, the total wage bill may exhibit very different patterns than the wages of the targeted workers. We control for log employment in these event studies, so the effects can be roughly interpreted as wages per worker and net out the effects of the overall growth in employment. However, this is a noisy measure of wages²⁵ and we cannot distinguish the trained workers from the untrained workers, new hires from incumbents, or even the targeted occupations from complementary ones.

²⁵Wages are the total wage bill paid out by the firm in a given quarter (summed over the year), while number of employees is measured at a snapshot date in the month. This measure will therefore be less accurate for firms experiencing strong churn or net growth in employment.

As we have already seen, training firms shift away from hiring in the higher-paying professional occupations in the years following training and lower skill requirements. These shifts could lower the average wage at the firm, drowning out any gains for the incumbent, trained workers. The dotted lines in Figure 8 add controls for the occupational mix of job postings, but this imperfect control for the composition of the stock of workers at the firm does not substantively change the effects on wages.²⁶

5.3 Effects by Market Concentration


Section 4 established that firms facing more competition in their local labor market are significantly more likely to receive grants to train their workers. Labor market conditions may also affect the impact of these grants. Firms facing a strong poaching threat may be particularly constrained from funding an optimal level of worker training in the absence of public subsidies, suggesting that the employment and vacancy composition effects of grants could be particularly strong for these firms. Greater competition and poaching risk may also force firms to pass a larger share of any productivity gains from training on to their workers by raising wages.

Figure 9 plots the effects of grant receipt on these two key outcomes separately for grant-receiving firms in local labor markets with an HHI above or below 0.15. This threshold is a focal point for Department of Justice merger guidelines, and product markets above this point have been considered “moderately concentrated.” The strong selection of training grants by market concentration illustrated in Section 4 constrains us from setting the threshold any higher; only 13% of firms receiving training grants operate in markets with an $HHI > 0.15$.

We find some evidence in support of the hypothesis that firms facing stronger competition shift employment and hiring more in response to grants. Five years post-grant, firms in low HHI markets are about 13% larger than pre-grant whereas firms in high HHI markets are about 6% larger. A z-test rejects that the effects are equal across groups with 99% confidence. We find less evidence that greater labor market competition drives firms to increase wages following training, though all the caveats for this imperfect measure of wages remain.

²⁶To include these controls, we must restrict the sample to firm-years with vacancy data from BG. The “BG Sample” plots in Figure 8 illustrate that this restriction alone does not change the qualitative story.

Figure 9: Event Studies by HHI



Notes: See figure 5 for regression specification information. HIGH HHI subsample includes grant-receiving firms located in markets with an $HHI > 0.15$ and the matched controls (regardless of market), while LOW HHI shows the complement. HHI calculated by commuting zone-industry, as discussed in section 4. Bars indicate 90% confidence intervals with standard errors clustered on matched pair.

5.4 Discussion

In this section, we have shown that establishments grow post grant both in terms of the number of employees and the number of vacancies. Training firms also shift vacancies away from higher-paying occupations and college requirements. Despite this compositional shift, average wages at the firm are unchanged. These effects either occur after the typical training window (1-2 years post receipt) or persist well after. We therefore interpret these effects as reflecting the changing nature of production after training is complete, rather than direct effects during the training window.

Grants appear to help firms shift to a long-term higher growth trajectory. Employment and job postings continue to grow for years following grant-supported training. Downskilling may follow from this growth as expansion tends to happen from the bottom (Engbom et al., 2023), consistent with our finding that post-grant job postings grow most in frontline sales and production roles. Demand for both targeted and non-targeted positions can grow with subsidized training, as long

as, for the latter case, the groups are production complements (Katz and Murphy, 1992; Restrepo, 2015). It may be that firms had a bottleneck in the production process and, once resolved, the firm is able to produce at scale and grow.

After training, firms grow disproportionately in areas that have fewer barriers to entry for low-skilled workers. That is an especially interesting result for policy makers, given that at baseline training firms appear to be good places to work (i.e., larger, higher wages, more established). When combined with the fact that grants are more likely to be used in more competitive labor markets, our results suggest that frictions in private sector human capital provision are prevalent but that public funds can induce firms to train in these markets.

6 Heterogeneity by Training Targets

For a subset of the states in our sample (California, Kentucky, Massachusetts, New Hampshire, and New Jersey), we have text descriptions of firms’ training plans taken from the grant applications. We use these descriptions to identify which broad occupation categories the training is directed towards. Because of the large number of training plans and their varied formats, we use Open AI’s Generated Pretrained Transformer (GPT) 3.5, a large language model (LLM), to classify each firm’s text. We tested multiple prompt formulations and found that directing the model to classify ads into four coarse occupation categories yielded more accuracy, measured against a human-coded test sample, than prompting the model to select more detailed occupations and then aggregating. See appendix D for detail.

Both conceptually and empirically, training plans can map into multiple categories. For instance, the example from the sign company (appendix figure E.3) proposes both production and office skills and specifically mentions cross-training employees to diversify from their current specializations. Our prompt to GPT allows the algorithm to identify multiple occupations associated with each training plan, with a probability weight assigned to each.

Table 2 reports the proportion of training plans that are categorized in each occupational grouping. We report two measures: Any Mention (col. 1) shows the proportion of plans in an occupational category allowing for a plan to be in multiple categories and Top Mention (col. 2-7) shows the proportion of plans in an occupational category based on the highest classification score. The most common types of training are in the professional skills group (60% any mention and 51% top mention), and in production skills (42% any mention and 29% top mention). These overall averages are somewhat distorted by Massachusetts, which provides nearly half of all the training

descriptions and awards disproportionately in professional skills.

Table 2: Proportion of Training Plans that Include Each Skill Group

	Any Mention	Top Mention					
		All	CA	KY	MA	NH	NJ
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Professional	0.660	0.505	0.409	0.143	0.676	0.533	0.281
Administrative/Sales	0.374	0.168	0.318	0.286	0.126	0.133	0.219
Service	0.120	0.036	0.046	0	0.0182	0	0.060
Production	0.416	0.291	0.227	0.571	0.180	0.333	0.440
Number of Grants	1290	1290	22	7	716	15	530


Notes. This table reports the proportion of training plans that were characterized as containing training in the four occupation groupings. For ‘Any Mention’, each plan can be in multiple categories, so the columns will not add to 1. For ‘Top Mention’ columns, categories are mutually exclusive and based on the category with the highest classification score. These statistics are calculated using the BG-only match data conducted outside the BLS enclave; see Appendix B for more details.

While many grants (37%) mention administrative or sales skills, less than 17% list it as the primary skill being trained for. Service related skills (usually customer service training) are even less commonly targeted. As such, we limit the analysis in this section to grants that target professional or production skills. We report results using the top mention categorization; using any mention produces similar results.

Figure 10 illustrates the words most commonly included in training descriptions we identify as targeting professional or production workers as the top mention. Grants targeting professional workers are strongly focused on leadership and management skills. Grants targeting production workers have a more dispersed focus, with common terms focusing on efficiency and lean manufacturing and some suggesting information technologies (“computer”, “software”).

We next explore heterogeneity in skill requirements in the BG-only sample as a function of training targets, both in terms of the baseline skill mix and the impact of training. For the former, we focus on the pre-training period and regress ad shares of skill requirements in firm-year cells on an indicator for whether the firm received a training grant in a sample of all control firms and firms receiving the indicated type of training. As above, we weight by number of ads in the cell. We do not observe what skills untreated firms would have trained in if they received grants. To capture some of the variation in skill needs across firms, we control for industry fixed effects in the pre-period comparisons. For measuring the impact of training, we focus on difference-in-difference

Figure 11: Skill Requirements by Target Occupation



Notes: The top panel plots differences in mean ad-level characteristics for firms receiving grants targeting Professional or Production workers vs. all control firms, controlling for industry fixed effects. The bottom panel plots difference-in-difference estimates separately by training type, controlling for firm and year fixed effects as well as a grant year indicator and a main effect for after (using placebo treatment years in the control group). All figures indicate 90% confidence bands when clustering on state. This figure uses the BG-only sample; see appendix B for details.

The upper left panel shows that firms targeting professional skills were much more likely to specify a wide range of skills in the pre-period: college and experience requirements, as well as keywords related to cognitive and computer skills. These skills are well-known to be highly related to professional jobs (Autor, 2019). Firms training in professional skills also had a higher share of postings in professional occupations (“prof occ”) compared to control firms and their ads scored higher on the professional skills index. These firms had outsized demand for workers in the target occupations and for the skills being trained for. This pattern is reassuring that the grant proposals do contain true information about firms’ labor demand needs.

The bottom left panel shows how skill demand evolves post-training. Firms that trained in professional skills see significant declines in demand for education, experience, cognitive skills, and computer skills. They also reduce their demand for professional occupations and skills related to professional jobs. Finally, they significantly shift their demand towards production-related occupa-

tions (“prod occ”) and skills. Before training, they had relatively less demand for these occupations relative to control firms in their same industry. Recall that manufacturing is disproportionately represented in grant-receiving firms, see figure 2, so even firms that target professional skills are likely to employ production workers.

For firms receiving training in production skills, the patterns are different. The upper right panel of figure 11 shows that they did not disproportionately demand production occupations or production related skills prior to receiving a training grant. If anything, they require more general skills (such as college and computer skills) and demand more jobs in professional occupations. Post-training, they continue their need for these general skills with significant increases in education and cognitive requirements and a noisy but large positive for computer skills. They also increasingly demand professional skills, though not professional occupations. Instead, they shift more towards production occupations and production skills. When we probed further (not shown), we found that these increased general and professional skill requirements are actually more pronounced within ads for production workers, rather than ads for occupations outside the targeted area. In sum, the firms that target production skills increase demand for production occupations post-training but also elevate skill requirements for these positions.

Table 3 explores difference-in-differences estimates in QCEW outcomes and the quantity of vacancies compared to their nearest-neighbor matches.²⁹ Firms training professional workers follow a pattern that is largely consistent with the aggregate results, consistent with these firms comprising a majority of all training firms. These firms posted substantially more BG ads before receiving grants and experience strong growth in log employment and number of BG job posts after grant receipt. This subset of professional-targeting firms also demonstrates significant increases in log wages per worker when controlling for shifts in the occupation mix of new hires. This control is important because, again consistent with the aggregate findings, this group of training firms decreases demand for a wide range of highly compensated skills and roles after receiving the grant.

Firms training production workers also experience meaningful, though smaller, increases in log employment and hiring after receiving the grants, relative to their matched control firms. These firms paid slightly more than their comparison firms before training, but—as in the aggregate results—do not show growth in wage bill per worker following grant receipt.

While we caution against inferring too much from the small sample of firms with grant descriptions, these patterns suggest two distinct use cases for training grants. In the first and most common case,

²⁹For these outcomes, which exhibited strong pre-trends in the aggregate results, we further reduce our sample by using the nearest neighbor matched control sample. For the ad characteristics explored above, we took advantage of the larger sample sizes and note that treated and control firms looked similar on pre-period BG outcomes.

Table 3: Wage and Growth Outcomes by Training Type

	Training Type			
	Professional		Production	
	Pre-period diff'l (1)	Diff-in-diff (2)	Pre-period diff'l (1)	Diff-in-diff (2)
Number of Posts	2.79	9.16*** (2.15)	0.54	3.03*** (0.63)
Observations		13629		7877
Log Employment	0.03	0.10*** (0.01)	0.02	0.05*** (0.01)
Observations		16776		11390
Log Wages	-0.06	0.04*** (0.01)	0.14	-0.02 (0.03)
Observations		3848		1473

Standard errors in parentheses clustered on nearest-neighbor match pair.

*** p<0.01, ** p<0.05, * p<0.1

Notes: Each panel summarizes a different dependent variable. Columns labeled (1) report differences in mean characteristic for firms receiving the indicated training type versus their nearest-neighbor matches. Columns labeled (2) report difference-in-difference estimates separately by training type, controlling for firm and year fixed effects as well as a grant year indicator and a main effect for after (using placebo treatment years in the control group). The difference-in-differences analysis also restricts to the nearest neighbor matched sample and clusters standard errors on match pair. The log wages outcome controls for log employment and the ad share distribution across broad occupation categories, analogous to the dashed line in figure 8, and therefore restricts to firm-year observations with non-zero BG posts.

growing firms may be blocked from further expansion by a shortage of managerial and operations related skills. Prior to receiving the grant, these firms posted many vacancies and demanded these kinds of professional skills that can become increasingly important during periods of rapid institutional change. Post-training, they grow even more rapidly and hire more, but are able to lower the skill requirements for new hires and no longer have outsized demand for professional positions. Controlling for these changes in the kinds of workers hired, these firms raise wages after grant receipt. These patterns are largely echoed, albeit sometimes more noisily, in the aggregate results in the previous section. On the whole, then, our evidence points to a plurality of training grants going to firms that need deeper managerial infrastructure to grow. Training resolves these needs and also results in growth in low barrier-to-entry positions, which may now be possible because of deeper skills at the top. Based on the wage evidence, it is likely that productivity improves as well.

In contrast, the smaller set of firms targeting production skills appear to be primarily pivoting their workforce. As illustrated in Figure 10, these firms propose training in production tasks, but with a focus on adaptability, efficiency, and technology. The pattern of changes for these firms is consistent with making the targeted workers more effective. Post-training, these firms hire more in the targeted occupations, but increase skill demands for these roles rather than lowering them. Many states highlight a desire to help firms and workers keep up with the pace of technological change, particularly in production tasks (see appendix section A). The focus on higher skill levels within production jobs aligns with management literature highlighting that today’s manufacturing workers will need digital skills and adaptability given the increasing adoption of new production technologies.³⁰ These firms emphasize hiring in cognitive and computer skills both pre and post training (in Figure 11), skills that are complementary with automation machinery (Hershbein and Kahn, 2018). It could be that training is intended to transition production workers to work alongside automation technology rather than displacing them, consistent with some evidence on technological change in Germany (Battisti et al., 2023).

7 Conclusions

Public-private incumbent worker training programs have the potential to improve outcomes, relative to typical public-sector training programs that tend to have disappointing results. Direct input by employers on the types of skills they need can align training with employment prospects. Public funding overcomes under-investment when employers worry that the workers they pay to train may be poached away and workers lack the resources or awareness to find training on their own.

In this paper, we compile a dataset of training grants that are allocated to private companies but administered by state governments using public funds. Exploiting unique linkages between the grants, the U.S. business registry, and the job postings of participating firms, we evaluate the characteristics of firms and markets that apply for and receive grants and then examine impacts of program participation. We find that grants are allocated to larger, older, faster growing firms that tend to hire more skilled workers. They are allocated to firms operating in labor markets that are larger and more competitive, with greater poaching risk. Finally, we find that grant participation facilitates growth that is disproportionately concentrated in lower skilled positions. Even as firms lower requirements for new hires, they keep payroll per worker steady.

Overall, our findings are inconsistent with place-based development motivations. In particular, we

³⁰Here are two examples of recent consulting reports: https://www.ey.com/en_us/industries/advanced-manufacturing-realized/prioritizing-next-generation-skills-for-manufacturing, <https://themanufacturinginstitute.org/new-report-dives-into-the-skills-needed-for-modern-manufacturing/>.

do not see grants allocated to small or under-developed markets or to firms that are new to the state but have a larger presence elsewhere. We do not see grants allocated to megafirms that might hold out-sized influence. Finally, we see grants having actual impact on labor inputs, ruling out perfect crowd out of private investments.

This collection of facts is consistent with the idea that training grants help resolve a market failure that prevented training from happening in the private market. After program participation, these high-quality firms reduce barriers to entry, either because they have learned they can train workers rather than imposing up front skill requirements or because training resolved a specific need for firms who can now grow in complementary jobs. The existence, allocation, and effects of these programs speak to a seminal literature in economics on the frictions associated with human capital provision in the private sector. By leveraging these unique programs, we highlight that Beckerian frictions are likely present in the private sector, and public funds can help to alleviate these barriers to training.

References

- Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association* 105(490), 493–505.
- Abadie, A. and G. W. Imbens (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica* 74(1), 235–267.
- Abadie, A. and J. Spiess (2022). Robust post-matching inference. *Journal of the American Statistical Association* 117(538), 983–995.
- Abraham, K. G. and M. S. Kearney (2020, September). Explaining the decline in the us employment-to-population ratio: A review of the evidence. *Journal of Economic Literature* 58(3), 585–643.
- Acemoglu, D. and J.-S. Pischke (1998, 02). Why Do Firms Train? Theory and Evidence*. *The Quarterly Journal of Economics* 113(1), 79–119.
- Almeida, R. and P. Carneiro (2009). The return to firm investments in human capital. *Labour economics* 16(1), 97–106.
- Altonji, J. G. (1993). The demand for and return to education when education outcomes are uncertain. *Journal of Labor Economics* 11(1, Part 1), 48–83.
- Andersson, F., H. J. Holzer, J. I. Lane, D. Rosenblum, and J. 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*.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021a). Synthetic difference-in-differences. *American Economic Review* 111(12), 4088–4118.
- Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021b). Synthetic difference-in-differences. *American Economic Review* 111(12), 4088–4118.
- Ashenfelter, O. (1978). Estimating the effect of training programs on earnings. *Review of Economics and Statistics* 60, 47–57.
- Ashenfelter, O. and D. Card (1985). Using the longitudinal structure of earnings to estimate the effect of training programs. *Review of Economics and Statistics* 67, 648–660.

- Autor, D. H. (2019, May). Work of the past, work of the future. *AEA Papers and Proceedings* 109, 1–32.
- Azar, J., I. Marinescu, M. Steinbaum, and B. Taska (2020). Concentration in us labor markets: Evidence from online vacancy data. *Labour Economics* 66, 101886.
- Banerjee, A. and E. Duflo (2004). Dp4681 do firms want to borrow more? testing credit constraints using a directed lending program. *Labour Economics* 15(I18), J31.
- Barron, J. M., D. A. Black, and M. A. Loewenstein (1987). Employer size: The implications for search, training, capital investment, starting wages, and wage growth. *Journal of Labor Economics* 5(1), 76–89.
- Bartik, T. J. (2017). A new panel database on business incentives for economic development offered by state and local governments in the United States. Working paper, Upjohn Institute.
- Battisti, M., C. Dustmann, and U. Schönberg (2023). Technological and organizational change and the careers of workers. *Journal of the European Economic Association* 21(4), 1551–1594.
- Becker, G. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. University of Chicago Press.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of Political Economy* 70(5, Part 2), 9–49.
- Belley, P. and L. Lochner (2007). The changing role of family income and ability in determining educational achievement. *Journal of Human capital* 1(1), 37–89.
- Berger, D., K. Herkenhoff, and S. Mongey (2022, April). Labor market power. *American Economic Review* 112(4), 1147–93.
- Black, D. A., L. Skipper, and J. A. Smith (2023). Firm training. *Handbook of the Economics of Education*, 287.
- Booth, A. L., M. Francesconi, and G. Zoega (2003). Unions, work-related training, and wages: Evidence for british men. *ILR Review* 57(1), 68–91.
- Brunello, G. and M. De Paola (2008). Training and economic density: Some evidence form italian provinces. *Labour Economics* 15(1), 118–140.
- Caliendo, M., D. A. Cobb-Clark, C. Obst, H. Seitz, and A. Uhlendorff (2022). Locus of control and investment in training. *Journal of Human Resources* 57(4), 1311–1349.

- California Panel Members (2022). The employment training panel 2021-24 strategic plan. https://etp.ca.gov/wp-content/uploads/sites/70/2022/04/ETP_2021-24_StrategicPlan-2022.pdf [Accessed: September 10, 2024].
- Card, D., J. Kluve, and A. Weber (2018). What works? a meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association* 16(3), 894–931.
- CareerSource Florida (2015). Qualified targeted industries for incentives. https://careersourceflorida.com/wp-content/uploads/2015/11/Targeted_Industries.pdf [Accessed: September 10, 2024].
- Clarke, D., D. Pailanir, S. Athey, and G. Imbens (2024). On synthetic difference-in-differences and related estimation methods in stata. *The Stata Journal* 24(4), 557–598.
- Commonwealth Corporation (2024). Workforce training fund program: General program guidelines. <https://commcorp.org/subprogram/wtffp-general-program-guidelines/> [Accessed: September 10, 2024].
- Crump, R. K., V. J. Hotz, G. W. Imbens, and O. A. Mitnik (2009). Dealing with limited overlap in estimation of average treatment effects. *Biometrika* 96(1), 187–199.
- Dalton, M. R., L. B. Kahn, and A. I. Mueller (2025). Do online job postings capture job vacancies? *Mimeo*.
- Dehejia, R. H. and S. Wahba (2002, 02). Propensity score-matching methods for nonexperimental causal studies. *The Review of Economics and Statistics* 84(1), 151–161.
- Deming, D. and L. Kahn (2018). Skill requirements across firms and labor markets. *Journal of Labor Economics* 36(S1), S337–S369.
- Dorn, D. (2009). *Essays on inequality, spatial interaction, and the demand for skills*. Ph. D. thesis, University of St. Gallen.
- Dustmann, C. and U. Schönberg (2009). Training and union wages. *The Review of Economics and Statistics* 91(2), 363–376.
- Engbom, N., C. Moser, and J. Sauermann (2023). Firm pay dynamics. *Journal of Econometrics* 233(2), 396–423.
- Forsythe, E., L. B. Kahn, F. Lange, and D. Wiczer (2022). Where have all the workers gone? recalls, retirements, and reallocation in the covid recovery. *Labour Economics* 78, 102251.

- Fuller, J. B., M. Raman, et al. (2017). Dismissed by degrees: How degree inflation is undermining us competitiveness and hurting america’s middle class. *Published by Accenture, Grads of Life, Harvard Business School*.
- Goldin, C. and L. Katz (2008). *The Race betwee Education and Technology*. Harvard University Press.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of econometrics* 225(2), 254–277.
- Hara, H. (2017). Minimum wage effects on firm-provided and worker-initiated training. *Labour Economics* 47, 149–162.
- Hashimoto, M. (1982). Minimum wage effects on training on the job. *The American Economic Review* 72(5), 1070–1087.
- Heckman, J., I. Hidehiko, J. A. Smith, and P. Todd (1998). Characterizing selection bias using experimental data. *Econometrica* 66, 1017–1098.
- Heinrich, C., P. Mueser, K.-S. Jeon, and D. Kahvecioglu (2013, 12). Do public employment and training programs work? *IZA Journal of Labor Economics* 2, 6.
- Hershbein, B. and L. B. Kahn (2018, July). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review* 108(7), 1737–72.
- Hirano, K., G. W. Imbens, and G. Ridder (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71(4), 1161–1189.
- Hollenbeck, K. (2008). Is there a role for public support of incumbent worker on-the-job training? Working Paper 08-138, Upjohn Institute.
- Hollenbeck, K. and W.-J. Huang (2015). Net impact and benefit-cost estimates of the workforce development system in washington state. Technical Report 13-029, Upjohn Institute.
- Holzer, H. J. (2023, November). Can workforce development help us reach full employment? Working Paper 16624, IZA.
- Holzer, H. J., R. N. Block, M. Cheatham, and J. H. Knott (1993). Are training subsidies for firms effective? the michigan experience. *ILR Review* 46(4), 625–636.
- Imbens, G. W. and Y. Xu (2024). Comparing experimental and nonexperimental methods: What lessons have we learned four decades after LaLonde (1986)? Working paper.

- Jones, D. C., P. Kalmi, and A. Kauhanen (2012). The effects of general and firm-specific training on wages and performance: Evidence from banking. *Oxford economic papers* 64(1), 151–175.
- Katz, E. and A. Ziderman (1989). *General training under asymmetric information*, Volume 170. World Bank Publications.
- Katz, L. F. and K. M. Murphy (1992). Changes in Relative Wages, 1963-1987: Supply and Demand Factors. *Quarterly Journal of Economics* 107(1), 35–78.
- Katz, L. F., J. Roth, R. Hendra, and K. Schaberg (2022). Why do sectoral employment programs work? Lessons from WorkAdvance. *Journal of Labor Economics* 40(S1), S249–S291.
- Kerr, W. R. and R. Nanda (2009). Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship. *Journal of Financial Economics* 94(1), 124–149.
- Konings, J. and S. Vanormelingen (2015). The impact of training on productivity and wages: firm-level evidence. *Review of Economics and Statistics* 97(2), 485–497.
- Laboissiere, M. and M. Mourshed (2017). Closing the skills gap: Creating workforce-development programs that work for everyone. *McKinsey & Company*.
- LaLonde, R. J. (1986). Evaluating the econometric evaluations of training programs with experimental data. *American Economic Review* 76(4), 604–20.
- Lynch, L. M. and S. E. Black (1998). Beyond the incidence of employer-provided training. *ILR Review* 52(1), 64–81.
- McCall, B., J. Smith, and C. Wunsch (2016). Chapter 9 - government-sponsored vocational education for adults. Volume 5 of *Handbook of the Economics of Education*, pp. 479–652. Elsevier.
- Michigan Department of Labor & Economic Opportunity (2024). Going pro talent fund program overview. <https://www.michigan.gov/leo/-/media/Project/Websites/leo/Documents/WD/Program-Overviews/Going-PRO-Talent-Fund.pdf> [Accessed: September 10, 2024].
- Mincer, J. A. et al. (1974). Schooling, experience, and earnings. *NBER Books*.
- Mohrenweiser, J., T. Zwick, and U. Backes-Gellner (2019). Poaching and firm-sponsored training. *British Journal of Industrial Relations* 57(1), 143–181.
- Muehlemann, S. and S. C. Wolter (2011). Firm-sponsored training and poaching externalities in regional labor markets. *Regional Science and Urban Economics* 41(6), 560–570.

- Murphy, K. M. and R. H. Topel (2016). Human capital investment, inequality, and economic growth. *Journal of Labor Economics* 34(S2), S99–S127.
- Negoita, M. and A. Goger (2024). State-level policies to incentivize workplace learning: Impacts of a california publicly funded employee training program. *Economic Development Quarterly* 38(1), 15–24.
- Neumark, D. and H. Simpson (2015). Place-based policies. In *Handbook of regional and urban economics*, Volume 5, pp. 1197–1287. Elsevier.
- O’Connell, S. D., L. F. Mation, J. Basto, and M. Dtuz (2019). Making public job training work: evidence from decentralized targeting using firm input. *mimeo*.
- Patnaik, A., J. Venator, M. Wiswall, and B. Zafar (2022). The role of heterogeneous risk preferences, discount rates, and earnings expectations in college major choice. *Journal of Econometrics* 231(1), 98–122.
- Restrepo, P. (2015). Skill mismatch and structural unemployment. Working paper.
- Rice, M., C. Griffiths, S. Little, J. McAloon, R. Meyer, A. Preston, and E. Slape (2005). California employment training panel: Annual report 2004-05. Accessed: 2021-10-25.
- Rosenbaum, P. R. and D. B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* 70(1), 41–55.
- Ruggles, S., R. Flood, Sarah abd Goeken, M. Schouweiler, and M. Sobek (2022). Integrated public use microdata series: Usa, version 12.0. *University of Minnesota, Minneapolis*.
- Rzepka, S. and M. Tamm (2016). Local employer competition and training of workers. *Applied Economics* 48(35), 3307–3321.
- Schubert, G., A. Stansbury, and B. Taska (2022). Employer concentration and outside options. *Mimeo*.
- Slattery, C. (2025). Bidding for firms: Subsidy competition in the united states. *Journal of Political Economy* 133(8), 2563–2614.
- Starr, E. P., J. Prescott, and N. D. Bishara (2021). Noncompete agreements in the us labor force. *The Journal of Law and Economics* 64(1), 53–84.
- Stevens, M. (1996). Transferable training and poaching externalities. In A. Booth and D. Snower (Eds.), *Acquiring Skills: Market Failures, their Symptoms and Policy Responses*, Chapter 2, pp. 19–20. Cambridge University Press.

- Stockinger, B. and T. Zwick (2017). Apprentice poaching in regional labor markets. *ZEW-Centre for European Economic Research Discussion Paper* (17-013).
- Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225(2), 175–199.
- Van Horn, C. E. and A. R. Fichtner (2003). An evaluation of state-subsidized, firm-based training: the workforce development partnership program. *International Journal of Manpower* 24(1), 97–111.
- Vermont Agency of Commerce & Community Development (2020). Vermont training program annual report fiscal year 2019. https://outside.vermont.gov/agency/ACCD/ACCD_Web_Docs/ED/VTP/AnnualReport_FY19.pdf [Accessed: September 10, 2024].
- West Virginia Economic Development (2012). The governor’s guaranteed workforce program: Annual report 2012. <https://westvirginia.gov/wp-content/uploads/2019/06/ggwfp-annual-report-2012.pdf> [Accessed: September 10, 2024].
- Yeh, C., C. Macaluso, and B. Hershbein (2022, July). Monopsony in the us labor market. *American Economic Review* 112(7), 2099–2138.
- Young, J., M. Ngatia, and E. W. Dillon (2026). Reevaluating causal estimation methods with data from a product release. *arXiv preprint arXiv:2601.11845*.
- Ziems, C., W. Held, O. Shaikh, J. Chen, Z. Zhang, and D. Yang (2024). Can large language models transform computational social science? *Computational Linguistics* 50(1), 237–291.

Appendix

A Policy Details on Training Programs

Breadth and Funding


We focus on public-private incumbent worker training programs, which are characterized by a firm-level proposal for state funds to share in the cost of training its own workers. At the national level, the Workforce Investment Act of 1998 allowed a small use of federal funds for such state-sponsored programs and this allocation was expanded in the Workforce Innovation and Opportunity Act of 2014 (WIOA). WIOA allows states to spend up to 20% of their allocated federal funds on incumbent worker training grants.

Beyond the federal level, state-level programs that provide funding for public-private training have existed since the 1960s with the majority of programs beginning in the 1980s and 1990s. In addition to WIOA funds, states use a combination of revenue from state unemployment taxes, general appropriation funds, and training-specific taxes to provide grants directly to firms to train incumbent or newly hired workers. A survey of 30 states by the Upjohn Institute in 2006 (Hollenbeck, 2008) found that states were investing around \$550 to \$800 million into public-private training partnerships, which is analogous to about 1% of what private firms spend on training. However, these programs have been largely overlooked by researchers since the WIOA expansion.

We conducted a comprehensive search of state incumbent worker training programs by browsing state training websites and combing program annual reports for detailed data. We tracked programs where the primary training grant recipient is an individual firm. Out of the fifty states and DC, we identified 45 which have programs that meet this criteria and 40 for which we can find aggregate annual expenditures on these programs. Figure A.1 reports average annual spending per-capita for these 40 states.

Of these 40 states, 18 have parsable firm-level data on program participation. Throughout all analyses, we restrict our attention to these 18 states, indicated in Figure A.1 with stars. The median spending among states with firm-level data is approximately \$2.60 per capita (Michigan), and the largest spender is New Mexico at approximately \$10 per capita.

Figure A.1: Per-Capita Spending on Public-Private Incumbent Worker Training Programs



Note: Stars indicate states which report firm-level data and are used in our analytical sample. Average per-capita spending on public-private training grants in author-collected data. We restrict attention to states that publish aggregate state-level data on spending. Per-capita spending is defined as the total dollars granted to firms in a state per fiscal year divided by the working age population (19 to 64 year olds) in that state with population data taken from the Current Population Survey (2013-2019). No data includes both states which have a program but do not report spending and states which do not have an identified program.

Stated Program Motivations

In promotional materials and program reports, most states reference a desire to improve the overall quality of jobs workers can attain and to target mismatches between worker skills and firm needs. For instance, Massachusetts asks applicants to “address selection criteria associated with job growth or increases in skills/opportunities of low-skill or low-wage workers” (Commonwealth Corporation, 2024). Similarly, Michigan hopes its program will “address skill shortages by reskilling and upskilling” (Michigan Department of Labor & Economic Opportunity, 2024).

Many states particularly highlight the challenges that both workers and firms face in keeping up with the pace of technological change. From Vermont’s 2019 program report, pg 3: “Advanced manufacturing is continuously evolving with more complex equipment which requires more technically advanced workers to program and maintain them. Meanwhile, employers continue to lose their content experts who are aging out of the workforce, often taking their institutional knowledge with them as few employers can afford succession planning. VTP is an excellent means of helping businesses to ‘up-skill’ existing employees allowing them to advance into the vacated positions” (Vermont Agency of Commerce & Community Development, 2020).

From a California report: “As rapid advancements in technology, automation, and artificial intelligence reshape the economy and the nature of work, more needs to be done to promote high-quality

jobs and economics security for workers, families, and communities”(California Panel Members, 2022).

Several states also indicate some place-based development goals. West Virginia describes their program as “play[ing] an important role in attracting new enterprises and encouraging the growth and expansion of the state’s existing companies”(West Virginia Economic Development, 2012).³¹ Half of the states in our sample list prioritized industries in their program descriptions. For example, California prioritizes manufacturing, healthcare, biotechnology, information technology, construction, agriculture, and logistics firms and in particular “targets firms threatened by out-of-state competition or who compete in the global economy”(Rice et al., 2005), while Florida targets “businesses able to locate in other states and serving multi-state and/or international markets”(CareerSource Florida, 2015).

Finally, states sometimes mention a desire to bolster economically disadvantaged labor markets, workers, and firms. Six states prioritize firms in areas with more disadvantaged workers. States often design their programs to ease the burden for smaller firms. For example, Maine requires firms with over 100 employees to pay 50% of training costs, firms with between 51 and 100 employees to pay 25% of training costs, and firms with less than 50 employees have zero required contribution. Ten of the eighteen states explicitly prioritize small businesses, with some states such as Michigan or Arizona providing additional points in their rubrics for businesses below a certain employee count.

Process

The 18 programs we study share some common features, but vary significantly in process, scope, and focus. Some states reserve training for incumbent workers while others require firms to hire new workers to train. In practice, 11 states allow for both incumbent and newly hired workers, 6 provide funding only for incumbents, and 1 limits to newly hired workers.

States vary in the extent to which they enforce their priorities for these programs through a competitive and discretionary review process. Six states evaluate grants using published scoring rubrics, although two of these states funded virtually all applications in practice during our sample. Michigan, which rejects a meaningful share of applications, has a 50-point rubric covering industry priorities, training provider quality, diversity considerations, post-training certification for workers, wages at the firm, and size of the funding request. Three other states do not publish a rubric, but

³¹Four states – New Hampshire, New Jersey, Oklahoma, and West Virginia – extend eligibility to firms that intend to physically relocate to the state, rather than only offering grants to firms already in the state. In contrast, Florida, Louisiana, and Ohio all require firms to have been located in the state for a minimum period of time before application.

describe a competitive review process in which not all applicants receive funding. In contrast, the other states in our sample do not mention any discretionary review of the training proposal in the application instructions and three states say explicitly that they will fund all applications on a first come, first served basis each fiscal period as long as program funds remain. States may effectively ration training grants by imposing high administrative burdens in the application and follow up reporting process. California, which funds the vast majority of applications that reach the final review board, has so many application requirements that firms typically hire expert consultants to navigate the process.

We document that 15 out of 18 states require grant-receiving firms to report employment status and wages of trained employees to the state.³² Some states use these reports to enforce requirements that firms retain workers who go through the training for a specified time post-training and/or pay them a specific wage. For example, Vermont requires that at the completion of training, the firm must pay a wage that equals or exceeds a ‘livable wage’ (\$15.33 as of 2022). Firms in Michigan must provide a company payroll query at three-months post-training reporting the name, hourly wage, hire date, and termination date (if applicable) for all employees trained, and they do not receive full reimbursement for training costs unless the trainee retained employment for 90 consecutive days post-training. These employment and compensation requirements will further influence which firms find it worthwhile to apply for training grants.

Many states structure the program to provide workers with credentials that can be carried across firms. Though some states allow for training to be internal (i.e., on-the-job), a number of states either explicitly require that training take place off-site through the state/community college system or a third party provider. Four states— Idaho, Indiana, Michigan, and Ohio— verify that workers have an industry-recognized credential at the end of training.³³

States put caps on the amount of funding the firm can apply for ranging from \$1,000 per worker in Idaho to \$8,000 per worker in Arizona.³⁴ Figure A.2 summarizes the distribution of grant dollars per worker, which is available for 75% of grants in our database. The median value is around \$1,100 dollars, though there is a sizable right tail so the mean (\$2,240) is considerably higher. Considering the typical training duration, these values amount to about \$20-\$40 per worker-week.


³²West Virginia also requires post-training reports from the firms, but information is not available on what these reports must include. There is no available information on whether New Hampshire or Oklahoma require post-training reports.

³³For example, firms in Ohio must provide the state with copies of a class roster, transcript, or a copy of the certificate for each trainee in order to receive reimbursement for the training. Maine’s program partners with the community college system, creating credit and non-credit based courses at specific colleges to meet the training needs of firms.

³⁴Some states cap total grant amount rather than per worker amounts. Grant size caps range from \$70,000 per grant in New Hampshire to \$850,000 per grant in California.

Employers cannot recoup much of their salary outlay. Instead, money can cover training materials and infrastructure, and small contributions for the opportunity cost of time. In most states, firms must provide some amount of matching funds (typically 50% of training costs).

Figure A.2: Grant Amount per Trainee



Note: Density plot of grant dollars per trainee across grants in author-collected data. Solid vertical line is the median; dashed line is the mean. For clarity, we omit from the figure (but not the mean and median calculations) grants with more than \$15,000 per trainee, 2% of our database.

B Matching Grants to QCEW and BG Data

We use firm name plus geography to match training participants to QCEW establishments, limiting attention to grants allocated from 2010-2019 (80% of our collected data). We first regularize employer names by removing common components such as LLC or “the”, removing punctuation, standardizing common word stems, etc. We then look for matches on exact (cleaned) name and county. When an exact match is not available, we use fuzzy matching techniques to find similar names across datasets, while relying on common geography to identify higher quality matches. Once training grants are matched to QCEW, we take advantage of the QCEW-BG match produced by Dalton et al. (2025) to bring in ad characteristics.³⁵ The resulting dataset uses firm name-county pairs as its unit of observation – the most detailed level at which we can match. Throughout, we refer to these name-county pairs as establishments or firms.

³⁵Note, for this latter match, we must restrict attention to the 70% of BG vacancy postings that specify an employer name. Ads with a missing name tend to be jobs posted by recruiting agencies.

Table B.1 summarizes the grants data and our matches to QCEW and BG. The full sample contains 13,375 cleaned grants averaging about \$92,000 in annual grant money. When available in the data, we observe that an average of 94 workers are to be trained, 58 of which are promised to be new hires. Average grant dollars per trainee is around \$2,000.

We are able to match 95% of the grants to an establishment QCEW. Columns 2 and 3 compare grant characteristics for matched versus unmatched grants. The grants that cannot be matched are larger in dollar amount and number of trainees. We also report the method used to match firms. The vast majority are matched on exact firm name after the initial clean, though we do pick up a non-trivial number of matches with the fuzzy match.

Of the QCEW matches, we are able to match 85% to a firm that posts at least one ad in BG. Columns 4 and 5 compare the BG matched to unmatched samples, among the QCEW matched grants. Again, grant dollar amounts are larger in the unmatched sample, while number of trainees and new hires is smaller. Dalton et al. (2025) as shown that small firms are less likely to post in BG. However, the grants that do not match to BG might be overall a noisier sample as indicated by their much lower exact-match rate to QCEW (48%, compared to 75% among the BG matched grants).

For a sub-set of our analyses, we re-run the linking algorithm on a BG-only dataset. This linkage allows us to conduct analyses for skill outcomes outside of the data enclave. Unlike the QCEW-BG-grant match, we cannot observe in this data set if a year where we observe no vacancies is due to the firm not existing or due to the firm not posting any vacancies online. For this reason, we prefer the QCEW-BG-grant data for analyses with number of vacancies as the outcome. However, for skill-outcomes which measure the proportion of ads that report a skill, the BG-grant match is equivalent as zero ad years are dropped from both samples.

Table B.1: Summary Statistics of Training Grants across Merge Samples

	(1)	(2)	(3)	(4)	(5)
	All Cleaned Grants	QCEW Match		BG Match	
		Matched	Unmatched	Matched	Unmatched
Grant Dollars	92137 (191109) N=13249	90837 (188184) N=12564	115976 (237387) N=685	87633 (167363) N=10650	108667 (276163) N=1914
# Trainees	94.39 (210.91) N= 9966	92.67 (206.33) N=9400	122.89 (274.82) N= 566	96.34 (214.73) N=7808	74.72 (157.67) N=1592
# New Hires	57.60 (1443.53) N= 5335	58.05 (1485.58) N= 5035	50.04 (130.32) N= 300	60.85 (1645.58) N=4092	45.91 (182.68) N=943
Grant Dollars per Trainee	2240.0 (4522.2) N= 9645	2259.8 (4600.7) N= 9107	1904.4 (2866.1) N= 538	2040.6 (3676.8) N=7600	3365.4 (7635.9) N=1507
Grant year	2015.2 (2.6)	2015.2 (2.6)	2015.3 (2.6)	2015.3 (2.5)	2014.8 (2.9)
Match to QCEW	0.95	1	0	1	1
Exact match	0.67	0.71	0	0.75	0.48
Match to BG	0.80	0.85	0	1	0
N (# Grants)	13375	12681	694	10750	1931

Notes: We report means of grant characteristics, as well as standard deviations in parentheses, and sample sizes (for the variables that are sometimes missing from the data). Grant data are assembled by the authors by reviewing state department of labor websites for training programs characterized by public funds flowing to individual firms to train their own workers. Column 1 includes the full sample of grants. Columns 2 and 3 compare grants that can be matched to the QCEW versus those that cannot, using the matching procedure described in the text. Columns 4 and 5 take the QCEW matched sample and compare grants that can be further matched to a firm in BG versus those that cannot, using the Dalton et al. (2025) merge, which follows a similar procedure.

C Nearest Neighbor Matching Algorithm

We match each grant-receiving firm without replacement to their one most similar untreated firm, considering only firms in the same two-digit industry within states with available training grant data for the reference year. To avoid capturing spillover effects within our control sample, we exclude all untreated firms in industry-county markets where at least 20% of workers were at firms

that received a training grant in any year within two years of the reference year. This exclusion removes control firms that are most likely to experience significant spillover effects while retaining a large and representative control sample. A stricter exclusion criteria that dropped all markets that ever received a grant in our data would remove virtually all large markets and limit our ability to find good matches for all treated firms. The “reference year” in this matching process is the year of first grant receipt for treated firms and the randomly assigned placebo year for untreated firms. Each untreated firm is therefore only eligible to be selected as a match in one, randomly assigned, year. This choice simplifies and speeds up the process of matching without replacement at the cost of reducing the pool of eligible matches in each year. In practice, the pool of untreated firms is so large that this restriction does not affect match quality.

Within the set of eligible firms, we select the single best match for each treated firm based on minimizing the Minkowski distance between pre-period outcomes. For the QCEW sample, we use log employment in periods $t - 1$ to $t - 5$ relative to the reference year and number of ads in Burning Glass in $t - 1$ to $t - 5$. We use the BallTree nearest neighbor matching algorithm, implemented in SciKitLearn, to match efficiently. For the BG-only sample, we run the match in Stata and use number of job posts in periods $t - 1$ to $t - 5$ for all years when we can observe five pre-periods and in observed pre-periods for all other firms.³⁶ Log employment and/or number of ads are effective summary measure of the size and growth trajectory of firms. For firms with missing log employment for some years of the pre-period, which largely reflect true zeros when the firm was not active, we fill in a value of -1,000, which is sufficient to ensure that we almost never match a firm with positive employment in some pre-period year to a firm with no employment in that year. For ad data, we replace missing values with zero in the QCEW sample in years where we believe the firm is in operation and with zeros in the BG data in years post-2010.

Finally, we drop firms from the matched analysis if we are unable to find a close match. In practice we drop matched pairs where the mean difference index is greater than the 99th percentile of distance. Results are stable to instead dropping the worst 5% or 10% of matches. Matching directly on these three core firm characteristics, industry, pre-treatment log employment, and pre-treatment hiring behavior, is sufficient to resolve the main violation of parallel trends in our full sample analysis: firms that apply for and receive training grants grow faster than the average firm in the years leading up to grant receipt.


One concern with matching on all pre-periods is that any post-period effects that we see are attributable to over-fitting and mean reversion. Specifically, if the control firm matched to the

³⁶Because the BG data starts in 2010, we cannot observe all five pre-periods for grants received prior to 2015 and thus match on the years it would have been possible to observe given sample truncation. Within this sample of firms that match to BG at some point, we categorize firms without BG activity in an in-sample year as posting zero ads.

treatment has idiosyncratically high employment in the pre-periods which makes it a good match to the treated firms (which on average have higher employment), it is possible that it will revert down towards the population mean in the post-period, resulting in an upwardly inflated treatment effect. We address this concern in three ways.

First, Figure C.1 plots the mean values of the two matched outcomes for the grant-receiving firms and the selected matched control firms. After matching and trimming, both treated and control firms show modest growth in log employment in the pre-period and stronger, parallel growth in number of posts. For both outcomes, the treated firms show a clear trend break after grant receipt while the control firms do not. This figure provides evidence that grants do indeed impact firms and that outcomes are driven by treated firms, with no evidence of mean reversion within the control group.

Figure C.1: Mean Outcomes for Treated and Matched Control Firms




Notes: We plot average outcomes by event time for firms receiving grants and the set of untreated firms identified as matched controls – see the text and appendix C for details on the matching procedure. Left panel shows log employment measures in the QCEW; right panel shows the number of postings. We limit to firm-year observations with non-zero employment and the right panel further limits to years of BG availability. Within these conditions, we impute a value of zero postings if the firm did not show up in BG in that year.

Second, we assess the validity of our matching assumption with a placebo exercise. In our BG-only

nearest neighbor match, we run each control firm selected as a match for a treated firm through the matching algorithm a second time, identifying a match to the matched firm. We then re-run our event study estimates giving the original matched sample a placebo treatment at time 0. If our algorithm is sensitive to overmatching, we might estimate spurious “treatment effects” even in this placebo control-to-control sample because of mean reversion. Instead, as illustrated by Appendix Figure C.2, this placebo exercise yields clear zero effects.

Figure C.2: Null Effects with Placebo Match




Notes: See figure 5 for regression specification details. Here we plot the baseline nearest neighbor match results alongside a placebo match in the BG-only data set. The placebo match takes the firms selected as match controls for the treated firms and repeats the matching algorithm to find a second set of matching control firms. See Appendix C for more details.

Finally, we test whether our results are robust to alternate matching methods to create our control group. We compare our results for the in-enclave nearest neighbor match, the out-of-enclave nearest neighbor match, and two alternative matching methods: synthetic difference-in-difference and inverse probability weighting. Figure C.3 shows the event studies for the outcome ‘Number of BG Posts’ for all four outcomes.³⁷ Panel A shows the in-enclave match results show in Figure 5. Panel B shows the out-of-enclave match. We see the out-of-enclave match has similarly flat pre-trends and grant receipt is associated with similar hiring demand post-grant in the out-of-enclave results.

³⁷We focus on number of posts in our sensitivity analyses due to limited access to government resources to run sensitivity analyses on the log employment outcome.

Figure C.3: Robustness to Alternative Match Specifications



Notes: This figure reports coefficients estimated using equation 3, our event study regression, for number of postings measured in BG. Panel A (top left) is for the QCEW-BG-grant NN match, Panel B (top right) if for the BG-grant NN match, Panel C (bottom left) is an inverse propensity score weighted regression, and Panel D (bottom right) is a synthetic difference-in-difference specification. We control for establishment fixed effects, year fixed effects, and event dummy main effects. We plot coefficients on event dummies interacted with treatment. We also report the 90% confidence interval for standard errors clustered at the match pair level.

Panel C of figure C.3 shows the results using inverse probability weighting (Hirano et al., 2003). We calculate the propensity to receive treatment by regressing an indicator for being a treated firm on average pre-period values of the following: number of BG posts, number of skills per post, proportion of ads requiring college, proportion of ads requiring any education, proportion of ads requiring any experience, proportion of ads requiring computer skills, proportion of ads requiring cognitive skills, and proportion of ads in each of the four occupation groupings. We trim the predicted propensity to have an overlapping distribution for treated and control and then calculate the weight for each firm as $W_i^{IPV} = D_i \times \frac{1}{p(X_i)} + (1 - D_i) \times \frac{1}{1-p(X_i)}$. We then re-run the regression weighted by this inverse probability. Again, results are similar magnitude to the primary specification.

Lastly, rather than finding a single firm to act as a control, we use the synthetic difference-in-difference method following Arkhangelsky et al. (2021b). This method combines features of syn-

thetic control methods and difference-in-difference in that it selects a weighted set of control units that minimize differences in trends in the pre-period. Panel D shows the results in a synthetic difference-in-difference design with bootstrapped standard errors, calculated using `sdid` function (Clarke et al., 2024). We restrict the sample to three periods pre- and post- grant receipt due to the strongly balanced sample requirements of `sdid`. While our estimates are much noisier in this reduced, strongly balanced panel, we see even larger point estimates of the post-grant effects.

These exercises suggest that our results are not particularly sensitive to the method of choosing a control group. We choose to use the nearest neighbor match for our primary specification because both synthetic difference-in-differences and inverse probability weighting have attributes that make analyses on BG skill demand outcomes infeasible. Synthetic difference-in-difference requires a strongly balanced panel, which restricts our sample to grants for which we can see all pre- and post- periods. While this can be done for number of posts where we fill in unobserved periods with zeros, this method cannot be used to look at any of the outcomes related to skills and occupational mix requested in ads as periods without ads are missing. Similarly, we cannot easily implement inverse probability weighting for regressions where the outcome is ‘Proportion of Ads with X’ as these regressions are already weighted by number of posts. In contrast, the nearest neighbor match can be used for all outcomes of interest.

D Categorization of Training Grant Descriptions

For five states in our sample, California, Kentucky, Massachusetts, New Hampshire, and New Jersey, we have text descriptions of the firm’s training plans taken from grant applications. To better understand what types of skills firms are using these grants to develop, we classify each training plan into one of four occupation groupings: (1) Professional, (2) Production, (3) Sales and Administrative Support, and (4) Service occupations. Because of the large number of training plans and the varied format of these plans, we use Open AI’s Generated Pretrained Transformer (GPT) 3.5, a large language model (LLM) to classify each firm’s text into these categories.

To construct predicted labels for each training plan text, we first supply a system-level prompt to GPT- 3.5. These system level instructions serve as a meta-prompt for the model and outline how the model should respond to subsequent user-level prompts. Figure D.1 shows the system-level prompt that we supplied to GPT for the classification task, and Figure E.3 provides an example of the training plan texts that are fed in as user-level prompts for classification. Specifically, we provided

in-depth details on the objectives of the task, what each groups consists of and their corresponding Bureau of Labor Statistics SOC codes, and in what manner the model should respond. To generate a prediction for each training plan text, then, we fed in each training plan text one at a time as user-level prompts to the model and collected responses. Finally, similar to Ziems et al. (2024), we set the temperature of the model to 0 to reduce the variance in GPT responses and create reproducible results as much as possible. We set all other model parameters to their default values.

We generated 5 GPT-classified samples for 3,540 training plans scraped from grant applications. We then constructed the predicted occupational targets for each training plan by taking the mode across the 5 samples. For example, if the set of occupational targets (in order) predicted by GPT-3.5 are (Professional, Production), (Professional), (Professional, Production), (Professional), (Professional, Production), the final predicted targeting would be (Professional, Production). In the case that GPT-3.5 did not have a majority prediction across the 5 samples (at least 3 of the predictions matching), those training plans were handlabeled. A similar approach is discussed in Ziems et al. (2023), where the authors average LLM responses over 5 different types of system-level prompts in order to generate predictions. In total, GPT-3.5 had complete consensus (all 5 predictions were the same) for 2,474 training plans, majority consensus (at least 3 predictions were the same) for 3,189 training plans, and did not reach consensus (and therefore required hand-labeling) for 64 training plans.

To give a concrete example, the firm depicted in Figure E.3 is a sign manufacturer which received a training grant in California. Based on the text used to classify this firm’s training plan, this firm is listed as professional, production, and sales and administrative support. While the company itself is a manufacturing firm and some of the training related to production skills such as safety procedures for crane usage or sign installation, many of the training skills listed include white-collar skills such as working with computer software like Microsoft Excel, improving HR skills, or negotiation skills.

Figure D.1: System Prompt to GPT-3.5

Assistant is an intelligent chatbot designed to help determine the occupational targeting of workforce development grants.

Each string of text that Assistant will receive is the training plan outlined by a company that is applying for a workforce development grant.

Each training plan is targeted to one or more occupation groupings.

Assistant's task is to determine which occupation group(s) the training plan is targeting given the training plan text by first determining which 2-Digit (major) SOC code (as provided by the Bureau of Labor Statistics) the plan is targeting and then aggregating into defined occupation groups defined below.

Here are the possible occupation groups, their descriptions, and their corresponding 2-digit SOC Codes, as provided by the Bureau of Labor Statistics (BLS):

- 1. Group: Professional, Description: Highly skilled white collar occupations. BLS SOC Codes: 11, 13, 15, 17, 19, 23, 27, 29.*
- 2. Group: Sales & Administrative Support, Description: Routine white collar positions such as sales and office support. BLS SOC Codes: 21, 25, 31, 41 (excluding occupations with minor SOC codes starting with 412, which are Retail Sales Workers), 43.*
- 3. Group: Service, Description: Positions such as servers and personal care jobs. BLS SOC Codes: 35, 37, 39, and occupations with minor SOC codes starting with 412 (Retail Sales Workers).*
- 4. Group: Production, Description: Blue collar jobs such as construction, production, and related occupations. BLS SOC Codes: 33, 45, 47, 49, 51, 53.*

In the case that there are multiple occupation groups that Assistant thinks the training plan is targeting, Assistant must rank their choices in order of most likely (first) to least likely (last).

Assistant's answer should be presented as such: Groups: (group choices) ; Reasons: (reasons). Note that the group choices should be listed FIRST in the area denoted "(group choices)" and the reasons should be listed in the area denoted "(reasons)".

The first answer in the group choice list must be the group that Assistant thinks the training plan is primarily targeting. The last answer in the group choice list must be the group that Assistant thinks the training plan is least likely to be targeting, but is still a focus of the plan itself. Respond ONLY with the group name, not the number. If Assistant does not think the Training Plan is targeting an occupation group, it should not include it in the list.

For example, if Assistant believes the training plan is targeting the Professional and Production groups with Professional being the most likely, Assistant's answer should be formatted as:


Groups: Professional, Production ; Reasons: Professional Reasons, Production Reasons.

Lastly, if Assistant does not think the training plan contains enough information to make a prediction, Assistant should simply return "ERROR: Not enough information contained in training text."

Notes: This figure shows the system-level prompt fed into GPT-3.5. This prompt is also referred to as the Aggregated prompt.


E Appendix Figures and Tables

Figure E.1: Availability of Grant Data by State and Year



Notes: We summarize the range of years for which grant data are available for a given state. Grant data are assembled by the authors by reviewing state department of labor websites for training programs characterized by public funds flowing to individual firms to train their own workers. We include data from any program that lists individual employer participants.

Figure E.2: Size and Number of Grants by State



Notes: We plot characteristics of grants by state for the matched sample of grants (see table B.1). We report unweighted means of the number of grants per year, grant dollars and number of trainees. The latter is unavailable in a small number of states. For these additional descriptives on posts, we use a sample created outside the BLS enclave that uses only BG information to match grants to firms. See appendix B for details.

Figure E.3: Example of Training Description

OVERVIEW

For over 60 years, The Company has manufactured electric and architectural signs. The Company's products are used primarily for brand identification and business location visibility. basic sign product is composed of either steel or aluminum and decorated to reflect the customer's name and/or logo. As a full-service sign company, provides services ranging from initial design concepts to detailed plans. The Company also provides fabrication, installation and maintenance of their products. customers include major hotels, property management companies, building owners, shopping centers, and general contractors.

uses Computer-Aided Design and Computer-Aided Manufacturing software integrated for design and fabrication. This technology provides with a competitive advantage. In order to meet growing customer demands and preserve its market share, seeks ETP funding to train employees at company sites in Oakland and Stockton.

Training Plan

All the proposed training is new content designed to supplement previous training. While some of the types and topics appear to be the same, the content has been updated. The training will be delivered by in-house trainers and vendors.

Business Skills: This training will be delivered to Contract Control, Project Coordinators, Sales, and Managers. Training will assist The Company as they manage growth and new project initiatives and implement ongoing business changes, such as reforms in HR processes to support growth. Expanding the skillsets of employees reinforces Company's commitment to creating a high performance workplace. Topics such as Estimating, Human Resources, and Effective Communications will be delivered.

Commercial Skills: Training will be offered to Production Staff and Installation Staff. This training will cross-train employees and diversify their current specializations so that employees have broader skillsets. Training will improve the ability of individual employees to perform more functions and services in order to boost overall productivity, improve safety, and gain specific competencies. Crane Operations, Electric Sign Installation, and Rigging are some examples of topics delivered. Driving related training does not include required licensing requirements. Some training topics will be delivered by vendors that offer certifications to demonstrate gained competencies such as forklift driving. Certifications generally add value to employees readiness to accept higher skilled higher paying positions.


Computer Skills: Training will be offered to Administrative Staff, Sales Staff, and Management Staff. Products like Gaant Charts and Microsoft Excel are being used by key contractors. Staff needs to be proficient and current on the newest software skills.

Manufacturing Skills: Training will be offered to Production Staff and Engineers. This training will help speed product fulfillment. New machinery including; mill saw, trimming, drills, vacuum, sander and spray gun were purchaed to keep pace with business changes. Training topics include; Tools, Structual engineering, Welding, and Certified Welding Inspector.

Continuous Improvement: Training will be offered to all staff to improve efficiency. Training topics include; Improving Sales Skills, and Negotiations. Sales Staff will receive Sales Skills Training which combines new product knowledge and customer relations. Construction Methodology will be given to Engineers to enable them to competitively bid and retain customers.


Notes: Example of a company-specific training plan that can be found in the state grant proposal documents. In this example, the text outlined in red is scraped, preprocessed, and fed into GPT-3.5 as a user-level prompt to determine its occupational targeting. This company may or may not appear in our merged analytic sample.

Figure E.4: Training Grants and Market Concentration: Robustness to Number and Size of Grant Outcomes




Notes: See figure 3. Here we plot bin scatters of total grant dollars or total number of grants (both including zeros) in a market on the concentration of vacancies. Markets are defined at the CZ-by-two-digit industry level and concentration is the HHI of job vacancies posted in the market

Figure E.5: Training Grants and Market Concentration using occupation-by-CZ Market Definition




Notes: See figure 3. We divide markets (CZ-by-three-digit SOC occupation pairs) into 20 equally-sized bins based on the HHI of job vacancies posted in the market (see equation 2). We then plot the share of markets that received any grants and average HHI within each bin.

Figure E.6: Training Grants and Market Concentration by State-Level Competitiveness




Notes: See figure 3. Here we split states into those with a rigorous scoring rubric and competitive selection process (left) and those with no apparent rigor (right).

Figure E.7: Log Employment Outcome: Comparing Full QCEW Sample to BG-Matched Sample



Notes: See figure 5 for regression specification details. The BG samples in this figure restrict to establishments that post at least one ad in the year. We also report the 90% confidence interval based on standard errors clustered by match pair.

Figure E.8: Ad Shares across Broad Occupation Groups (BG)



Notes: See figure 5 for regression specification information. Here we restrict to establishments that had BG postings in the year. Outcomes include the fraction of postings in each of four mutually exclusive and exhaustive occupation groups: professional, routine white collar (Admin), low-skill service (service), and blue collar (production), see footnote 13). Bars indicate 90% confidence intervals with standard errors clustered on matched pair. This figure uses the BG-only sample; see appendix B for details.

Table E.1: Summary Statistics for Grant and Non-Grant Firms

	All Firms		NN Matched	
	Grant	Non-Grant	Grant	Non-Grant
Panel A: QCEW Sample				
Characteristic				
Employment	227.666	22.051	186.0	152.514
Wage per Worker	62100.829	43412.017	61584.319	57112.497
Age	16.925	13.497	17.078	17.085
Annual Growth Rate	0.076	0.058	0.07	0.05
# Semi-Firms	8667.0	1742145.0	8298.0	8298.0
Panel B: BG-QCEW Matched Sample: Semi-Firm Level				
Characteristic				
Employment	258.122	40.237	206.627	184.795
Wage per Worker	63349.115	49556.976	62690.828	60611.397
Age	17.303	14.086	17.435	17.804
Annual Growth Rate	0.078	0.064	0.073	0.048
# Semi-Firms	7042.0	619190.0	6733.0	5522.0
# BG Postings in t-1	23.76	4.43	13.45	13.40
Panel C: BG-Only Matched Sample: Ad-weighted				
Characteristic				
Education Req.	0.676	0.539	0.652	0.603
College Req.	0.452	0.285	0.435	0.353
Experience Req.	0.552	0.483	0.559	0.525
Computer Skills	0.330	0.267	0.343	0.307
Cognitive Req.	0.477	0.354	0.475	0.398
Professional	0.668	0.465	0.632	0.524
Admin	0.150	0.264	0.142	0.218
Service	0.0336	0.114	0.0397	0.0827
Production	0.0974	0.116	0.136	0.132
# of Ads	379664	2752852	224033	178649

Notes: This table reports characteristics for grant and non-grant recipients. Panel A uses the entire QCEW matched sample, while panel B restrict to the BG-QCEW matched sample and panel C to the BG-only sample (see Appendix B). Establishment-level characteristics and the number of BG postings are measured in the year before grant receipt; the employment growth rate measures the change between t-2 and t-1. BG ad characteristics are ad-weighted averages for the entire pre-grant period. For comparison, non-grant recipients are assigned a placebo grant year at random, excluding the first and last 2 years of operation. The occupation variables divide SOC occupations codes into four mutually exclusive and exhaustive groups: Professional includes SOC 11-19, 23, 27, 29; Sales/Admin is 21, 25, 31, 41 (excluding 412), 43; Low-skill Services is 35-39, 412; and Blue Collar is the remainder (33, 45-53).

Table E.2: Training firms and market characteristics: 3-digit Occupation-by-CZ


Dependent Variable	Any Grants Received (mean = 0.207)			
	(1)	(2)	(3)	(4)
HHI	-1.219*** (0.342)	-1.405*** (0.327)	-1.217*** (0.357)	-1.175*** (0.332)
HHI ²	1.928** (0.673)	2.231*** (0.735)	1.975** (0.701)	1.868** (0.648)
CZ unemp rate	-0.199 (0.590)	-0.584 (0.405)	-0.147 (0.568)	-0.274 (0.715)
New Market	-0.055 (0.035)	-0.087*** (0.027)	-0.046 (0.037)	-0.065 (0.046)
Employment (1,000s)	0.065*** (0.005)	0.059*** (0.005)	0.067*** (0.005)	0.066*** (0.006)
Wage (\$100s)	0.273*** (0.057)	0.293*** (0.085)	0.080 (0.176)	0.160 (0.124)
Emp growth			-0.049 (0.053)	-0.042 (0.049)
Wage growth			0.037 (0.065)	0.047 (0.062)
Leave-out State Emp			-0.009 (0.013)	
Leave-out State Wage			0.211 (0.207)	
Leave-out State Emp Growth			-0.012 (0.115)	
Leave-out State Wage Growth			0.115 (0.075)	
Leave-out Region Emp				-0.003 (0.003)
Leave-out Region Wage				0.027 (0.034)
Leave-out Region Emp Growth				-0.078 (0.052)
Leave-out Region Wage Growth				-0.039 (0.168)
State-Yr, CZ-Yr, Occ. FE	X	X	X	X
Occ-Yr, Occ-State		X		
Observations	77,224	77,224	77,224	77,224
R2	0.276	0.332	0.276	0.276

Standard errors in parentheses clustered on state

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Observations are 3-digit occ-by-CZ-by-year; standard errors are clustered on state. HHI, Employment, and Wages are occupation-by-CZ averages from 2010-12. Emp and wage growth are the rate of change in 2012 from 2010 for the occ-by-CZ. The CZ unemployment rate varies by occupation and year. The Leave-out State and Region variables are also at the occupation-by-geography level, averaged over 2010-12 or the rate of change over that period and leave out the focal CZ or state, respectively. Regression observations restricted to 2013-2019. Covariates are defined for the 13,902 markets that posted at least 50 ads in the baseline 2010-12 period and have coverage in the ACS, and other markets are considered “New”.

Figure E.9: Skill Requirements (BG) Event Studies



Notes: See figure 5 for regression specification information. Outcomes are the proportion of ads specifying the indicated skill requirement. Cognitive and computer skill indices are taken from Deming and Kahn (2018). Regressions with dotted lines control for the composition of ads across the 4 occupations groups in the year. Bars indicate 90% confidence intervals with standard errors clustered on matched pair. This figure uses the BG-only sample; see appendix B for details.

Table E.3: Robustness to alternative measures of market concentration

Dependent Variable	Received Grant (Mean: .2)					
	Occupation Mkt		Industry Mkt			
	(1)	(2)	(3)	(4)	(5)	(6)
HHI	-1.687*** (0.426)		-0.991*** (0.111)			
HHI ²	2.763*** (0.803)		1.036*** (0.132)			
Share of ads to top 3 Firms		-0.108*** (0.017)		-0.300*** (0.031)		
Share of Emp in top 3 Firms					-0.185*** (0.027)	
Industry Tightness						0.008*** (0.002)
Employment (1,000s)	0.075*** (0.005)	0.084*** (0.007)	0.043*** (0.009)	0.045*** (0.009)	0.049*** (0.010)	0.043*** (0.010)
Wage (\$100s)	0.262*** (0.061)	0.380*** (0.057)	0.804*** (0.227)	0.789*** (0.214)	0.696*** (0.227)	0.464* (0.228)
Observations	77,224	77,224	19,788	19,788	19,773	18,754
R-squared	0.217	0.219	0.274	0.279	0.265	0.260
Two-way FEs	X	X	X	X	X	X

Standard errors in parentheses, clustered by state
*** p<0.01, ** p<0.05, * p<0.1

Notes: See tables 1 and E.2 for regressions specification; FEs include state-year and industry. Industry tightness is the number of jobs posted in BG, averaged over 2010-2012, divided by 10,000 times the number of unemployment workers who previously worked in the industry as measured in the ACS in the same time period.

Table E.4: Event Study Coefficients: Log Employment, # Posts, Log Wages (QCEW)

	Log Emp.	# Posts	Log Wages	Log Wages	Log Wages
t-5	-0.006 (0.004)	-0.420 (0.638)	-0.009** (0.005)	0.012 (0.011)	0.012 (0.011)
t-4	-0.005 (0.004)	-0.203 (0.539)	-0.007* (0.004)	0.000 (0.008)	0.001 (0.008)
t-3	-0.009*** (0.003)	-0.617 (0.458)	0.002 (0.004)	0.013* (0.007)	0.013* (0.007)
t-2	-0.005* (0.003)	-0.174 (0.407)	0.002 (0.003)	0.007 (0.006)	0.007 (0.006)
t-1	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
t=0	0.054*** (0.003)	1.804*** (0.500)	0.004 (0.003)	-0.001 (0.005)	-0.001 (0.005)
t+1	0.080*** (0.005)	2.950*** (0.691)	0.004 (0.004)	-0.004 (0.006)	-0.004 (0.006)
t+2	0.101*** (0.006)	4.811*** (0.867)	-0.004 (0.004)	-0.011 (0.007)	-0.011 (0.007)
t+3	0.113*** (0.007)	6.034*** (0.997)	-0.005 (0.004)	-0.006 (0.008)	-0.006 (0.008)
t+4	0.128*** (0.008)	10.423*** (2.112)	-0.006 (0.005)	0.001 (0.009)	0.001 (0.009)
t+5	0.124*** (0.009)	9.690*** (2.120)	-0.012** (0.005)	-0.011 (0.010)	-0.011 (0.010)
Firm FE	X	X	X	X	X
Year FE	X	X	X	X	X
BG Sample				X	X
Occ Control					X
R2	0.964	0.533	0.989	0.991	0.991
Firm-Year Obs.	161689	139120	161687	38736	38736

Standard errors in parentheses clustered on matched pair

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports event study coefficients for regression specification 3 with log employment as the outcome in Column 1, the number of BG posts in Column 2, and log wages in Column 3-5. For log wages, column 3 is the nearest neighbor matched sample; column 4 restricts to observations that post in BG in the year; and column 5 adds controls for the proportion of ads across the 4 occupation groups in the year. These correspond to the nearest neighbor event studies in Figure 5 and Figure 8.

Table E.5: Event Study Coefficients: Composition of Vacancies in BG-Only Sample

Dep Var:	Occupation Ad Share				(5) Avg. Occ. Wage	Skill Requirement Ad Share	
	(1) Prof	(2) Admin	(3) Serv	(4) Prod		(6) College	(7) College w/ Occ Control
-5	-0.009 (0.016)	0.013 (0.012)	0.004 (0.006)	-0.001 (0.008)	-0.215 (0.386)	-0.005 (0.020)	-0.030 (0.018)
-4	-0.023* (0.013)	0.020* (0.010)	0.001 (0.005)	0.002 (0.007)	-0.513 (0.353)	-0.022 (0.017)	-0.038** (0.016)
-3	-0.008 (0.010)	0.017* (0.009)	-0.004 (0.006)	0.003 (0.005)	-0.054 (0.250)	-0.031*** (0.011)	-0.048*** (0.014)
-2	-0.003 (0.008)	0.005 (0.006)	-0.002 (0.003)	0.004 (0.005)	0.036 (0.203)	-0.010 (0.010)	-0.020 (0.013)
-1	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)	0.000 (.)
t=0	-0.012 (0.009)	0.009 (0.006)	-0.003 (0.003)	0.000 (0.005)	-0.221 (0.191)	-0.005 (0.010)	-0.001 (0.012)
1	-0.020** (0.010)	0.014* (0.007)	0.001 (0.004)	-0.002 (0.008)	-0.540** (0.226)	-0.018 (0.013)	-0.002 (0.013)
2	-0.037*** (0.010)	0.022** (0.010)	-0.001 (0.004)	0.006 (0.006)	-0.923*** (0.232)	-0.022 (0.015)	-0.014 (0.013)
3	-0.043*** (0.012)	0.021** (0.010)	0.003 (0.006)	0.011* (0.006)	-0.954*** (0.280)	-0.048*** (0.018)	-0.008 (0.013)
4	-0.072*** (0.027)	0.042* (0.025)	0.010 (0.006)	0.007 (0.007)	-1.477*** (0.549)	-0.075*** (0.029)	-0.011 (0.014)
5	-0.066* (0.036)	0.029* (0.017)	0.003 (0.009)	0.028 (0.022)	-1.296 (0.869)	-0.071* (0.037)	-0.009 (0.014)
Firm FE, Yr FE	X	X	X	X	X	X	X
Occ Control							X
Observations	1,468,820	1,468,820	1,468,820	1,468,820	1,468,820	1,468,820	54,156 R-squared
0.845	0.753	0.914	0.842	0.797	0.830	0.523	

Standard errors in parentheses clustered on matched pair

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports event study coefficients for regression specification 3. Outcomes are the share of ads across each broad occupation group (cols 1-4), ad-weighted average wage of the occupational mix (col 5), and share of ads requiring college without controls (col 6) and with controls for occupational shifts (col 7). All are in the nearest neighbor matched BG-only sample, include firm and year fixed effects and cluster standard errors by matched pair. These correspond to the nearest neighbor event studies in Figures 6 and 7.