Do Workforce Development Programs Bridge the Skills Gap?*

PRELIMINARY AND INCOMPLETE; PLEASE DO NOT CIRCULATE

Eleanor W. Dillon
Microsoft Research NE

Lisa B. Kahn
University of Rochester, NBER, and IZA

Joanna Venator
Boston College

Michael Dalton†
Bureau of Labor Statistics

September 20, 2024

Abstract

Most U.S. states have a workforce development program that offers firms grants to train their own workers. These training programs may help close skills gaps or may primarily serve local development goals. This paper explores the determinants and consequences of such programs. We create unique data linkages between participating firms and the Bureau of Labor Statistics business registry, as well as the Burning Glass job vacancy data (the near-universe of online job postings). We find that training grants are more prevalent in markets where firms face greater employee poaching risk, as well as in larger and higher-paying firms and markets. Using an event study and nearest-neighbor matching research design, we find that after training, firms experience growth in the number of postings and employees. Growth in job postings is concentrated in lower-skilled, front-line occupations and, even conditional on occupation mix, firms relax skill requirements after receiving a training grant. As such, program participation facilitates access to relatively high-quality firms. These low-skilled positions may complement those that received training or participating firms may have learned how to train workers themselves, rather than imposing up-front requirements. This collection of facts is consistent with the notion that these programs help overcome a market failure in updating worker skills.

*We thank seminar participants at Arizona State University, Carnegie Mellon, the Cleveland Fed, Columbia, the Markets with Frictions Tinos Edition, SITE: The Micro and Macro of Labor Markets, Society of Labor Economics, SEML, University of Michigan, and the University of Rochester. This project was funded in part by generous support from the Smith Richardson Foundation.

†E-mail: Eleanor.Dillon@microsoft.com; lbkahn@gmail.com; venator@bc.edu; 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 has only been 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). However, 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 cover the cost of the training themselves. 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 existing or newly hired) workers. Nonetheless, there are few existing 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 job vacancy data (BG).[2] To better understand the rationale for these programs, 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 program participation impacts labor demand, we then examine the impact of program participation on employment and vacancies using an event study and nearest-neighbor matching design.

[1]See for instance a recent 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.

[2]The 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.

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, and some states mention focusing on out-of-state competition, especially in under-developed markets. In practice, we find that grants are much more likely to be used in competitive labor markets, as measured by the concentration of firms hiring in the market or market tightness. We also find that grants are concentrated in larger and higher paying firms and labor markets, and 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 at new or young markets, to firms that are new to the state but not new overall, or to megafirms that might have outsized political influence.

Next, we analyze the impact of program participation on firm hiring and employment outcomes. We use an event-study design to compare treated firms to, first, all other firms and then a matched sample of similar firms.After grant receipt, vacancies and realized employment levels at participating firms increase relative to the control groups. Using job vacancies as a proxy for the composition of employment growth, we find that jobs shift away from professional occupations and toward lower-skilled, front-line positions. Employers also reduce requirements for education and related work experience post-training. These effects accrue over time and are unlikely to be driven solely by direct effects during the training period itself. Training may have helped to resolve a bottleneck in production so that firms are now able to operate at optimal scale and grow in complementary jobs. Or, these firms may have invested in "training capital" such that they are now willing to take a chance on less skilled workers.

The evidence we present suggests these grants resolve a skills gap which previously prevented the firm from operating at optimal scale. The fact that labor inputs change post grant receipt means that these grants are not simply 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. Furthermore, such firms are, on average, located in more competitive labor markets. This finding is consistent with the view that public-private incumbent worker training programs help to solve a Beckerian friction in which firms are reluctant to pay to train their own workers due to poaching risk, as opposed to a place-based incentive policy. Overall, these programs appear to accomplish exactly what they say they will: helping to increase access to high quality firms and potentially narrowing the skills gap. Importantly, our firm-level results show that public-private partnerships work to increase firm hiring of low-skill front-line positions, suggesting that these programs are an effective policy solution to address the gap in demand for those with limited formal education.

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 seminal 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 skills training in areas of local need and especially skills with greater poaching risk 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 evaluating the impacts of programs in Massachusetts (Hollenbeck, 2008), Michigan (Holzer et al., 1993), New Jersey (Van Horn and Fichtner, 2003), and Rhode Island (Angell et al., 2021). 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.

Our analyses uniquely 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 firms or employees are surveyed (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; 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 exploring how 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

[3] See 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.

[4] O'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.

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 contributing to a seminal and largely theoretical literature in labor economics on human capital (Becker, 1964; Acemoglu and Pischke, 1998)).[5]

While college graduates learn general analytical skills that help them shift tasks with changing skill demands, a large group of workers with limited 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 some light on the constraints that prevent firms from providing training without public support.

This paper proceeds as follows. The next section 2 provides institutional detail on the training programs we study, discusses motivations for these programs, and lays out empirical tests to disentangle which motivators are most important in practice. Section 3 describes data sources and summarizes characteristics of training firms. Section 4 relates training allocation decisions to market-level characteristics. Section 5 examines changes in employment and vacancies as a function of grant recipiency. Section 6 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. 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

[5] A number of papers have explored how poaching risk correlates with training provision in the European market, finding mixed support (Muehlemann 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.

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.

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

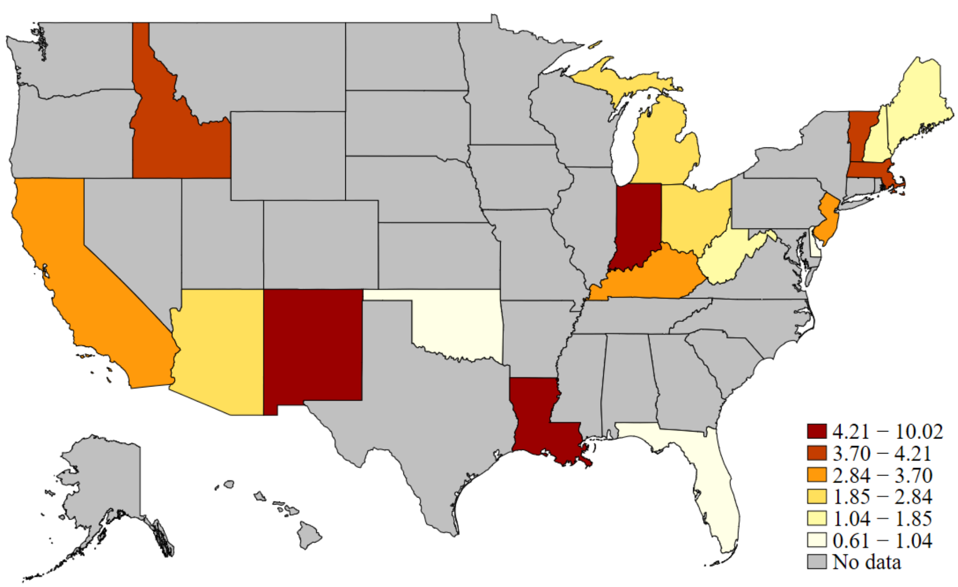

Note: Average per-capita spending on public-private training grants in author-collected data. We restrict attention to states that publish employer-level data. Per-capita spending is defined as the total dollars granted to firms in a state per fiscal year divided by the working age population (25 to 64 year olds) in that state with population data taken from the Current Population Survey (2013-2019).

We conducted a comprehensive survey of state incumbent worker training programs by browsing state training websites and combing program annual reports for detailed data. We track programs where the primary training grant recipient is an individual firm – rather than a worker or business consortium – to distinguish from traditional worker training programs. Out of the fifty states and DC, we identify 42 which have programs that meet this criteria.

Of these 42 states, 18 have parsable firm-level data on program participation. Throughout all analyses, we restrict our attention to these 18 states, indicated in Figure 1 with shading for different levels of average annual spending per-capita. The median spending is approximately \$2.60 per capita (Michigan), and the largest spender is New Mexico at approximately \$10 per capita. we describe the data we collect on these programs in more detail in the following section.

2.2 Program Administration

The 18 programs we study share some common features, but vary significantly in process, scope, and focus. In all states, the 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.[6] Length of training varies by state, 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.

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.

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).[7] Half the states in our sample list prioritized industries in their program descriptions. For example,

[6]We 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 (as in other programs that incentivize hiring the currently unemployed). 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.

[7]Four 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.

California "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, including California and Louisiana, prioritize firms in areas with more disadvantaged workers. States often design their programs to ease the burden for smaller firms. 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.

Process

States vary in the total administrative burden of applying for these grants. While some states report high rejection rates or describe a competitive process, others either have much less information on how they allocate grants or expressly state a first-come, first-serve approach. Six states evaluate grants using published scoring rubrics. For example, Michigan's 50-point rubric covers industry priorities, training provider quality, diversity considerations, post-training certification for workers, wages at the firm, and size of the funding request. California approves the vast majority of applications that reach the final review board, but publishes a long list of priorities and requirements up front, and firms typically hire expert consultants to navigate the process.

Most programs either give higher priority to firms which promise to increase wages following training or explicitly require that workers receive a particular wage. For example, Vermont requires that at the completion of training, the firm must pay a wage that equals or exceeds a 'livable wage' which is currently set as $15.33. We document that 15 out of 18 states require firms to report employment status and wages of trained employees to the state.[8] For example, 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.

In addition, 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

[8] West 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.

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.[9]

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.[10] Figure 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 per worker, 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. 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 2: Grant Amount per Trainee

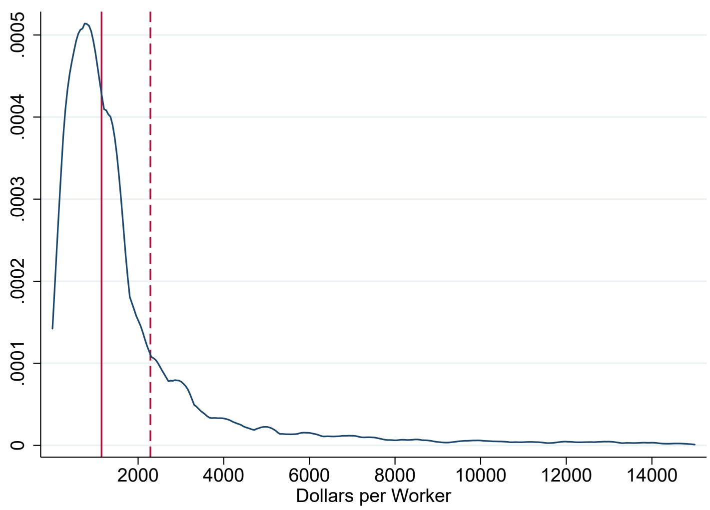

Note: Density plot of grant dollars per trainee across grants in author-collected data. Solid vertical line is the median; dashed line is the mean. For clarity, we omit from the figure (but not the mean and median calculations) grants with more than $15,000 per trainee, 2% of our database.

Between the limited dollar values, credentialing and pay raise requirements, and administrative overhead surrounding these grant programs, we expect substantial self-selection of firms. Firms will likely only apply when they can usefully train a large group of workers and/or meet the administrative hurdles of application and compliance. Therefore, it is not a priori clear that the

[9] For 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.

[10] Some 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.

ultimate recipients of these grants are always in line with the states' goals.

2.3 Conceptual Motivations for Empirical Analysis

We conclude this section by summarizing various economic theories that can explain why firms apply to these programs despite the administrative burdens and why states might provide them (regardless of their stated motivations). We lay out observable predictions of each theory to motivate our empirical analysis, which will attempt to identify which models appear most relevant in practice.

The canonical theories of human capital investment suggest that employers and employees who have already reached a work agreement should also be able to come to an agreement 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)). Small and young firms may also behave like individual workers as they face some of the same borrowing constraints as individual workers (Banerjee and Duflo (2004), Kerr and Nanda (2009)).[11]

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 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.

Economic literature also motivates the broader place-based development goals of the state. There is

[11] Minimum 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.

ample evidence that states use these 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, 2020).

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. These firms might also have out-sized political influence and be therefore better able to direct funds.

No matter the states' motivations in funding these programs, there is always a risk that public dollars will crowd out private investment. In the absence of any of the frictions discussed above, they could perfectly crowd out private dollars, in which case grants would have no impact on firm outcomes. If the grant money tips some firms over the margin of training an additional worker, we might see the impacts of such training on other labor inputs of the firm. There may or may not be many firms exactly on that margin. When grants are targeted toward areas where frictions in the provision of skills might arise, we expect there will be more of these marginal firms and we will consequently observe larger effects on firm performance.

Our analyses will proceed in two steps. First, we will describe the distribution of grant participants in terms of firm characteristics and labor market features. In light of the likely strong self-selection of firm applicants, it will be interesting to see whether grant allocations are consistent with stated place-based development goals. Furthermore, we will explore whether grants tend to be used in more competitive labor markets. A greater need for public funds when firms face poaching risk is consistent with under-investment in general skills due to market frictions.

Second, we will evaluate whether grant recipients change labor inputs following program participation, relative to an observational control group. It may impact overall growth if production was lexicographic in the skill being trained for. Also, once training is acquired, firms may shift demand from skill areas covered in the grant to complementary skills. Evidence of impacts implies that public funds are not simply crowding out private investment dollars. Rather, funds are being allocated to firms on the margin of training some number of workers. For any such firms, the firm-specific benefits would not outweigh their private training costs. However, combined with the analysis of how grants are distributed, we can inform whether government dollars are going towards areas where the social benefits to training outweigh the costs due to frictions in the provision of human capital.

3 Data

3.1 Hand-collected program data

After combing state websites and reports 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. The earliest year of data we collect is 2002 for California and the latest year of data we collect is 2019 for twelve out of the eighteen states. Appendix figure A.1 provides further details on the availability of grant data by year. In addition to firm name, the majority of in-sample states also report the county of participating firm, number of trainees requested, and value of the grant. Appendix figure A.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. We use these descriptions to identify which broad occupation category the training is directed towards. 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.[12] We also use these categories below when measuring firms' recruiting behavior. 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 into these categories. See appendix B for detail.

Table 1 reports the proportion of training plans that are categorized in each occupational grouping. Both conceptually and empirically, training plans can map into multiple categories. For instance, the training plan in appendix figure A.8 is for Arrow Sign Company, a firm in California that manufactures electronic signs, and proposes training in machinery as well as a range of basic office

[12]This grouping maps SOC occupation codes into four mutually exclusive and exhaustive SOC occupation code groups: Professional includes SOC 11-19, 23, 27, 29; 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).

skills. As such, the columns in table 1 do not sum to 1. The most common types of training are in the 'professional' skills group (59%), which can be thought of as skills used in high-skill, white collar occupations, and in 'production' skills (45%) which can be thought of as blue collar occupations such as construction or manufacturing jobs. These overall averages are somewhat distorted by the two states that provide the most training descriptions, Massachusetts and New Jersey. Massachusetts disproportionately awards grants to firms requesting training in professional skills (81% of grants) whereas New Jersey's grants are fairly evenly split across production, professional, and administrative/sales (i.e., low-skill white college occupations). Across the non-Massachusetts grants, there are fairly even proportions across production (49.7%) and professional skills (49.8%).

Table 1: Proportion of Training Plans that Include Each Skill Group

	All	CA	KY	MA	NH	NJ
Professional	0.592	0.674	0.364	0.812	0.530	0.488
Admin/Sales	0.392	0.389	0.545	0.352	0.220	0.419
Service	0.183	0.200	0.0909	0.0760	0.0900	0.238
Production	0.446	0.632	0.636	0.324	0.430	0.493
Number of Grants	2863	95	11	855	100	1802

Notes. This table reports the proportion of training plans that were characterized as containing training in the four occupation groupings. Each plan can be in multiple categories, so the columns will not add to 1.

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. This comprehensive administrative database provides our ground truth for select firm characteristics and also for the firm's survival each year.

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 labor 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. (2023) 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.[13]

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. (2023). 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 A.

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 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.[14]

[13] See 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.

[14] Most 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

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 8,495 grant firms and 1.5 million control firms. Appendix table A.2 provides summary statistics for grant firms (column 1) and this full set of control firms (column 2). 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 treatment and control group in the distribution of firm characteristics in Figures 3 and 4. These figures take the share of grant receiving establishments of a given characteristics (for instance, size bin or industry) and subtract the non-grant recipient group share.

Figure 3: Treatment-Control Differential in Distribution of Establishment Characteristics

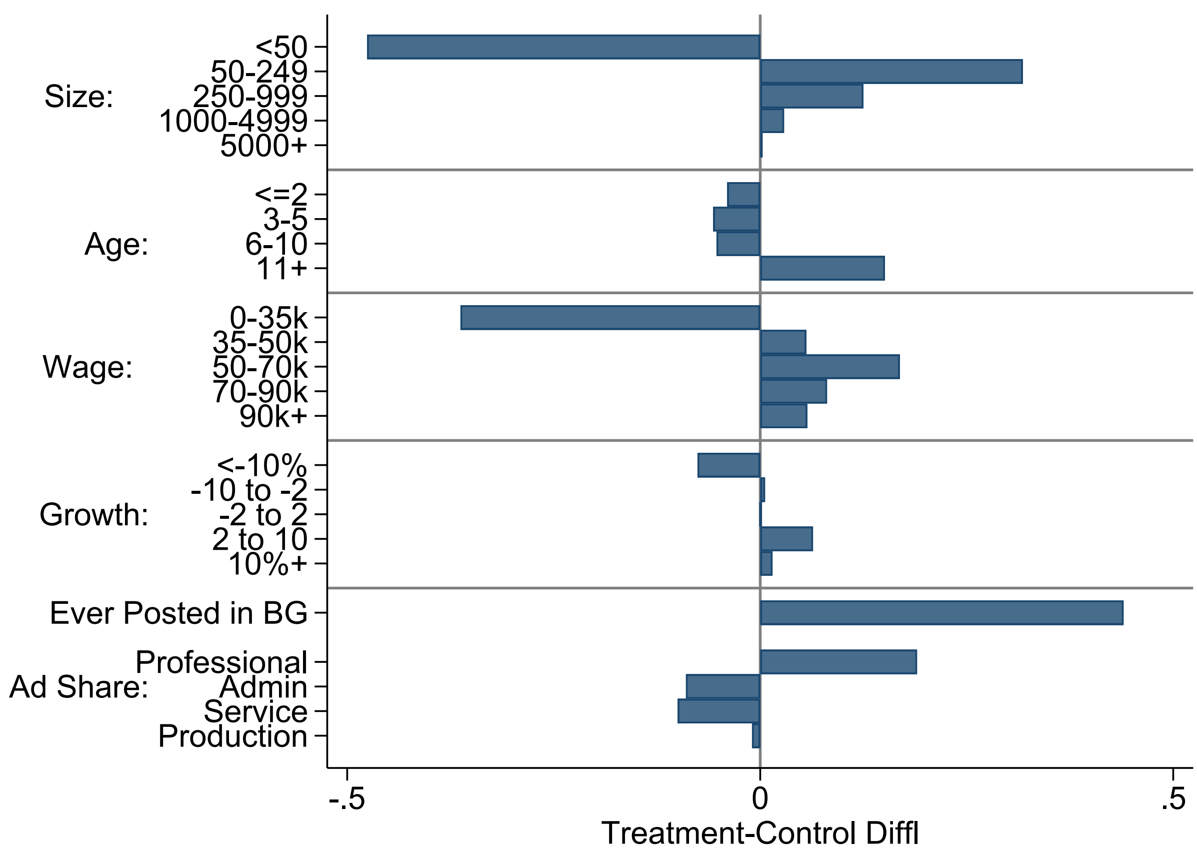

Notes: We plot the difference between the fraction of treatment establishments in a bin and the fraction of control establishments. We do this for characteristics in the year prior to grant receipt (or placebo year). Growth rate is measured as the t-2 to t-1 change in employment. The ad share across occupations in BG restricts to firms that post ads in t-1. See footnote 12 for definitions of the broad occupation categories.

firms, we exclude those with no more than 1 employee for average monthly employment, as this group of firms is highly unusual but represents a non-trivial fraction of establishments.

Beginning with the employment size distribution, 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). These distributional differences make grant firms larger on average, with recipients averaging roughly 200 workers, compared to the control average of about 25 workers. Interestingly, while grant firms are larger on average, we do not see an overrepresentation among megafirms – those with 5000+ employees.

Turning to firm age, we see that grant recipient establishments are older by about 3 years on average, with substantial overrepresentation (15 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 next look at wages. The wage concept in the data is measured 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. However, grant firms' payroll per worker is quite a bit higher, averaging about $20K more per worker. When we look at representation across the average wage distribution, grant recipients are substantially less likely to be found in the lowest wage bin (36 ppts) and much more likely to be found in the middle and high wage categories.

For growth rate, which we define as the percent change in employment between t-2 and t-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 – 82% can be matched to BG at any point, compared to only 38% 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 – averaging 41 per year, compared to 16. Panel B of appendix table A.2 also shows that, within the BG sample, differences in establishment characteristics across grant and non-grant recipients are similar to those in panel A.

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. Within the BG sample, grant firms specify skill requirements at higher rates: they specify an education (experience) requirement in 71% (60%) of ads, compared to 55% (47%) 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. Among the firms who use online hiring services, treated firms have a greater proportion of job ads asking for professional skills (65% relative to 46%) and are less likely to be searching for skills relevant to administrative/sales, service, or production occupations. Though not shown, these differences across grant and non-grant groups persist even after controlling for the differences in industries shown above.

Finally, grant recipients are concentrated in different industries than non-recipients. From figure 4, 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).

Figure 4: Industry Distribution across Grant and Non-Grant Recipients

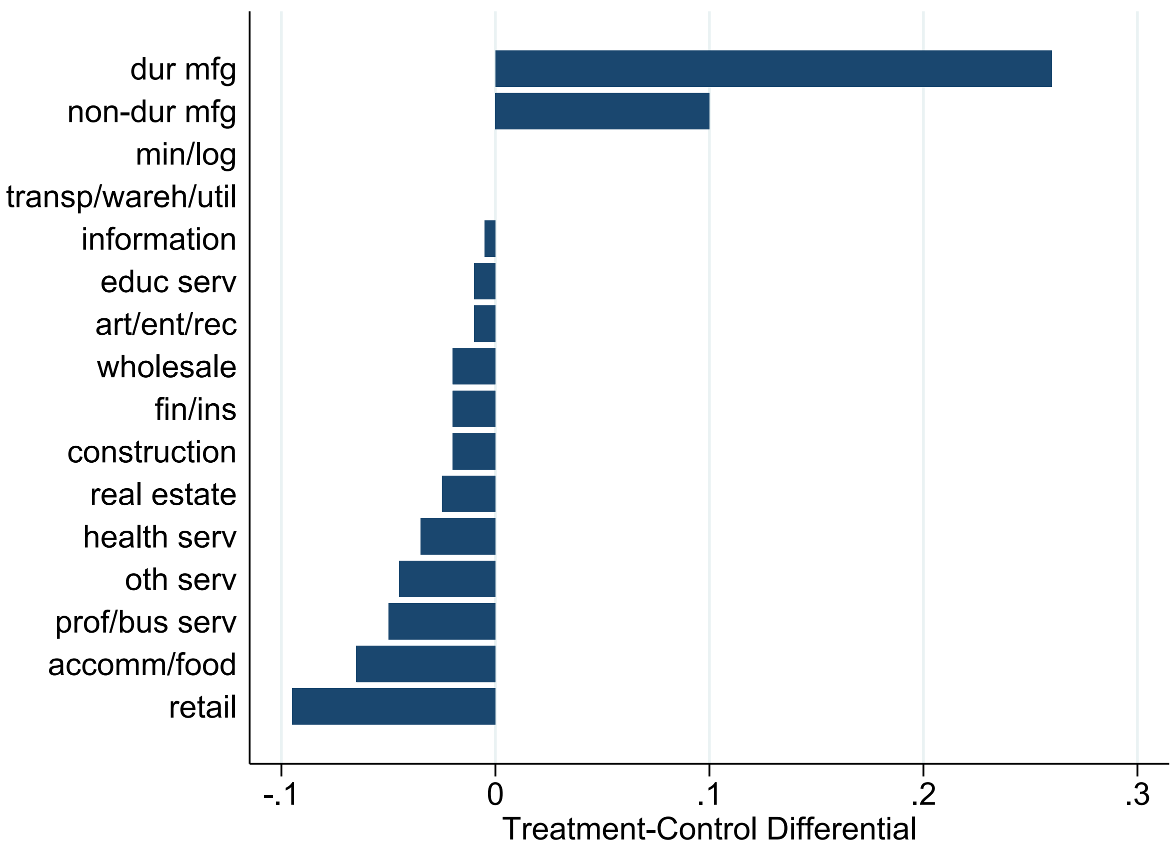

Notes: We plot the difference across treatment and control group in the share of establishments in two-digit NAICS sectors.

4 Grant recipiency 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 places 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, 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. Markets are defined by commuting zone, c, and skill, j, which we classify by either occupation or industry. Our baseline specification controls for state-by-year fixed effects ($\theta_{s(c),t}$), to examine the relationship between economic activity and grant allocation within the specific grant cycle, and skill (θ_j) fixed effects. We cluster standard errors by state to account for persistent state-level correlations in grant allocation decisions.

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

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.

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

The HHI in market cj is the sum of squared ad shares across all firms, k, posting in the market. A higher value on this index indicates that a greater proportion of job vacancies in a given market are from a small number of firms (i.e., a less competitive market).

Our preferred measure of poaching risk uses this vacancy-based market concentration index, particularly salient for labor demand, measured at the three-digit occupation-commuting zone level. 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. Occupations seem the most intuitive way to define these markets. However, grants are allocated to firms, not occupations, so we must impose an additional step to map firm-level grants to occupations. Using the ad distribution of the establishment, we allocate grants to the modal occupation among the firm's job postings. We explore both alternate measures of concentration and alternate definitions of markets, particularly by industry, and will show our results are quite consistent across all variations.

We also explore the relationship between grant receipt and a range of other market characteristics ($\mathbf{X}_{cj}$) such as size, average wage, employment growth, and health relative to neighboring markets. For the occupation-CZ-level analyses, we use American Community Survey data for 2010-2012 (Ruggles et al., 2022), combined with crosswalks between public-use micro areas from Dorn (2009), to calculate the average number of 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.[15] We also use the ACS to measure CZ-wide unemployment rates.

We take the mean level of annual employment and earnings in the market over the three year period, as well as calculating employment and wage growth rates between 2010 and 2012. To better understand economic activity in neighboring markets, we also calculate "leave-out" versions of these measures at the state-occupation or state-industry level (omitting the focal CZ-skill market from that calculation) and the Census division-occupation or division-industry level (omitting the focal state from that calculation).

All of the measures of market-level economic activity are calculated as the mean of the measure

[15] 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.

across the years 2010-2012.[16] We restrict our regression analyses of the propensity to receive a grant to the years 2013-2019, conditional on observing grant allocations in each state-year. 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 activity. 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 the measures of economic activity, to dummy them out.

4.2 Results

Figure 5 provides a bin scatter of the likelihood that the market receives at least one grant on the y-axis and the market-level HHI on the x-axis. We see that markets with 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. Appendix figure A.4 shows a similar relationship with 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 2, 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, marklet-level employment and wages, CZ-year unemployment, and two sets of fixed effects. State-by-year controls adjust for anything happening during the state's grant allocation decisions or the firm's application decisions that might be correlated with the key covariates – for instance, if more firms choose to apply when state budgets grow; occupation fixed

[16]We choose these years because they are the earliest years for which we have consecutive coverage of the Burning Glass data.

Figure 5: Training Grants and Market Concentration

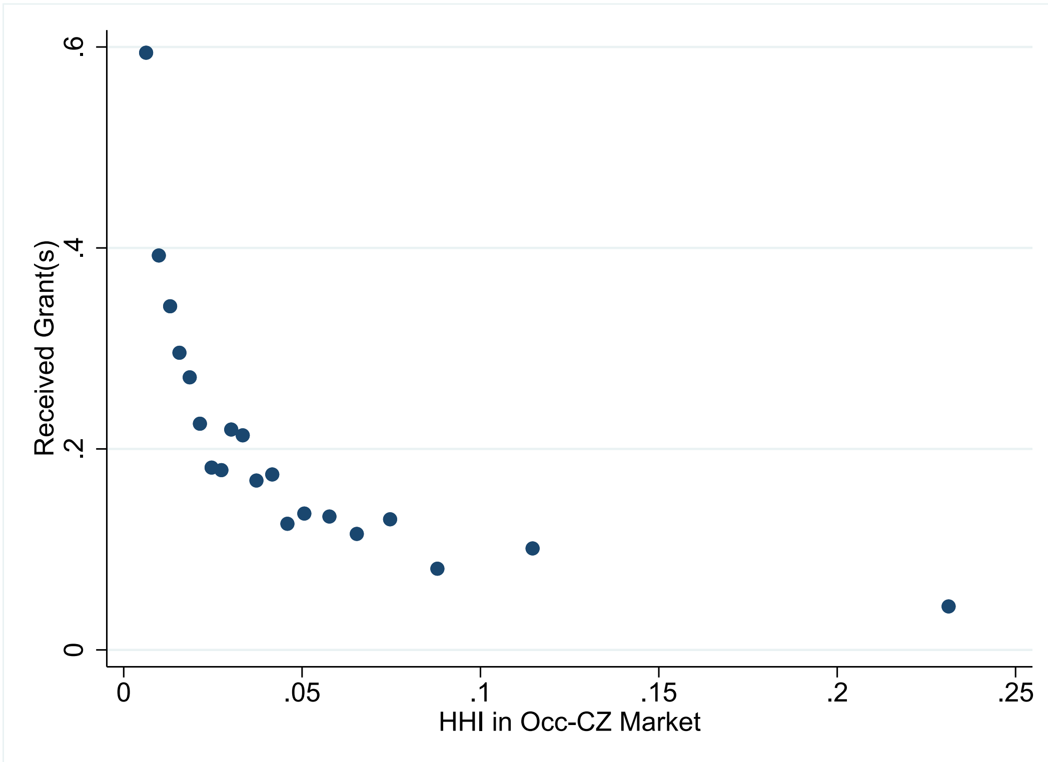

Notes: We divide markets (CZ-by-three-digit occupation pairs) into 20 equally-sized bins based on the HHI of job vacancies posted in the market (see equation 2). We then plot the average HHI and the share of markets that received any grants within each bin.

effects control for the possibility that certain occupations are in favor with state grant agencies and these also happen to be in more or less concentrated markets. To provide some context for the magnitude of the relationship, the mean and standard deviation of the HHI are 0.051 and 0.057, respectively, and the average market receives a grant with 20.6% likelihood, meaning that a one standard deviation increase in HHI from the mean is associated with a 6 ppt (30%) 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 an 8 ppt (40%) higher likelihood that the market received at least one grant; markets with a 1 ppt higher unemployment rate are slightly less likely to receive a grant (by about a third of a point), though the difference is not statistically significant.

Columns 3 and 4 test whether grant receipt is associated with the economic characteristics of

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

Dependent Variable	Any Grants Received (mean = 0.207)			
	(1)	(2)	(3)	(4)
HHI	-1.651***	-1.916***	-1.599***	-1.623***
	(0.417)	(0.390)	(0.424)	(0.407)
HHI^2	2.718***	3.212***	2.629***	2.653***
	(0.788)	(0.835)	(0.777)	(0.753)
CZ unemp rate	-0.368	-0.493	-0.288	-0.471
	(0.638)	(0.604)	(0.648)	(0.787)
New Market	-0.137**	-0.152**	-0.135**	-0.152**
	(0.052)	(0.054)	(0.055)	(0.064)
Employment (1,000s)	0.075***	0.072***	0.076***	0.075***
	(0.004)	(0.004)	(0.005)	(0.006)
Wage ($100s)	0.245***	0.301***	0.307	0.120
	(0.066)	(0.083)	(0.215)	(0.140)
Emp growth	-0.084	-0.061	-0.083	-0.083
	(0.071)	(0.075)	(0.078)	(0.071)
Wage growth	0.060	0.075	0.046	0.062
	(0.082)	(0.080)	(0.083)	(0.082)
Leave-out State Emp			-0.025*	
			(0.013)	
Leave-out State Wage			-0.064	
			(0.254)	
Leave-out State Emp Growth			0.016	
			(0.155)	
Leave-out State Wage Growth			0.130	
			(0.078)	
Leave-out Region Emp				-0.006
				(0.004)
Leave-out Region Wage				0.037
				(0.040)
Leave-out Region Emp Growth				-0.049
				(0.068)
Leave-out Region Wage Growth				-0.071
				(0.225)
Observations	77,224	77,224	77,224	77,224
R-squared	0.217	0.270	0.218	0.218
Occ, State-by-Year FEs	X	X	X	X
Occ-by-year,Occ-by-State		X		

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

Notes: Observations are 3-digit occ-by-CZ-by-year. 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 year. The State and Region variables are also at the occupation-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 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”.

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 occupation-CZ's in the state, the "leave-out state" market, while column 4 uses the population-weighted average of all other state-occupations in the census division, the "leave-out region" market. We revert to the original sets of fixed effects controls since these neighboring market variables have little or no variation within occupation-year or occupation-state. We find little evidence that characteristics of the occupation within the state as a whole impact the empirical grant distribution, nor do neighboring states. No coefficient is statistically significant at conventional levels, and confidence intervals are such that we can rule out fairly sizeable effects. The signs on the coefficients point towards grants in markets whose neighbors in the state are smaller and lower paying, with potentially faster employment growth. Though, again, these patterns are noisy and do not hold for out-of-state neighbors.

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. Figure 6 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. We do not see any evidence that indicates strategic choices by states to target new markets.

These results are robust to a range of alternative approaches. Column 2 of Table 2 illustrates robustness to the inclusion of occupation-by-year and occupation-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 programming jobs may tend to be located in concentrated markets. Occupation-by-state effects help control for the possibility that preferences for a given occupation 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 and at near double the magnitude.

Appendix table A.3 shows the results hold when defining markets by CZ and industry (at the 2-digit NAICS level) rather than occupation. While less aligned with our concept of market, industry has the benefit of a more straightforward mapping between firm-level grants and markets and allows us to include more precise controls for other market characteristics. HHI is still calculated with

Figure 6: Training Grants and # Ads Posted in the Market

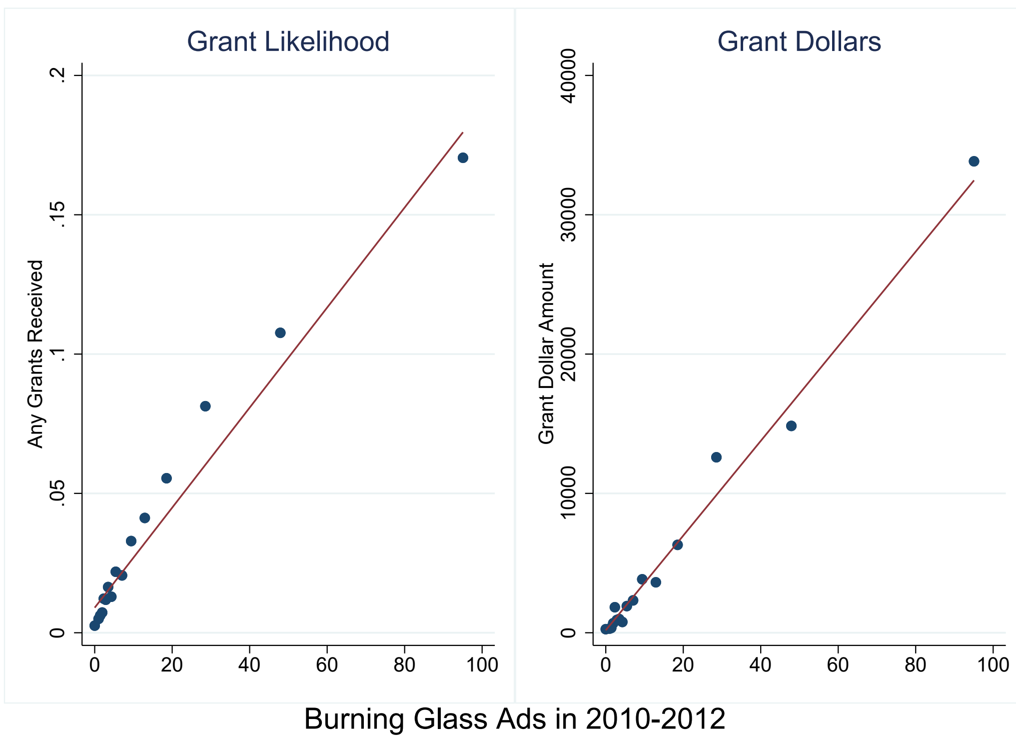

Notes: We divide markets (CZ-by-three-digit occupation pairs) into 20 equally-sized bins based on the number of ads posted in the market from 2010-12. We then plot the share of markets that receive at least one grant (left) or the average dollar amount per grant (right, including zeros) on the average number of ads in the baseline period. We restrict to markets that post no more than 169 ads (the 95th percentile) for visual clarity for the smaller markets, though the slope of the line is fairly similar when we include them.

Burning Glass, but we can use the full employment distribution from County Business Patterns to capture the other variables, rather than just the ACS survey. The more precisely measured local labor market variables possibly help reduce standard errors such that we see statistically significant relationships between neighboring characteristics and grant allocation. Reinforcing the suggestive finding at the occupation level, grants are significantly more likely to be allocated to markets whose in-state neighbors are smaller and faster growing.

In Appendix table A.3, we also control for the proportion of firms per market with less than 10 employees and the proportion of firms per market with more than 500 employees. If states target credit constrained firms, we would expect grant receipt to be positively associated with the proportion of small firms and negatively associated with the proportion of large firms. However, if anything, it looks as though grants are allocated to markets with a higher fraction of employment in larger firms.

Finally, we explore other measures of employment concentration, rather than the concentration of job vacancies. Results reported in appendix table A.4 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. 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

In summary, we find a strong and robust negative relationship between market concentration and grant allocation. This finding is consistent with the hypothesis that firms are reluctant to pay to train workers when they are competing heavily for talent within the market. In these instances, public sector subsidies can help solve market failures and, interestingly, we find training grants are much more likely to show up in these markets. It is not clear from these results whether this pattern is driven by state governments targeting these 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 competitive evaluation process

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

that results in only some applicants receiving grants) or firm-led (i.e., 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 A.3 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.

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). In fact, 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.

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 of an establishment, T, defined as the first year we observe the firm receiving any grant. We again cluster standard errors by state, the level at which treatment is determined. 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 issues that may arise with staggered adoption of treatment (Sun and Abraham, 2021). The event time indicators of interest are all interacted with the ever treated indicator – $\mathbb{1}(grant_i)$.

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

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 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.

We explore a wide range of outcome variables to better understand patterns in employment, wages, and vacancies. 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. In most states, we do not observe the specific skills that training is for, so our outcomes target understanding average changes in firm inputs.

Our identification strategy relies on the standard parallel trend assumption of a two-way fixed effects model: in the absence of the grant, the treated establishment's employment and vacancies would have followed the average trend to other establishments. This is, of course, a strong assumption. We might think that the type of firm which chooses to apply for and is awarded a grant is on a different growth trajectory than firms which do not apply for or receive a grant.

To address selection issues, we use several approaches. Our preferred controls include 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. In robustness checks, we add controls for industry-year to capture common sectoral shocks. When we include industry-year fixed effects in additional specifications reported in Appendix Section D, results are not substantively different though more noisily estimated.

We observed in our descriptive analyses that treatment and control firms differ on a wide range of observable characteristics. In particular, we see that grant recipient firms are growing at a faster rate prior to receiving a grant, suggesting a violation of the parallel trends assumption. As an alternative, we use a nearest neighbor matching design to find a control group that has similar trends to the treated group on a key set of characteristics. Specifically, we use a Ball tree algorithm nearest-neighbor search(Rosenbaum and Rubin, 1983; Abadie and Imbens, 2006) to identify a matched firm in the same industry that minimizes the 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 indicators for whether the firm posted in BG. 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 treatment or placebo treatment year. Following Abadie and Spiess (2022), we match without replacement, which allows us to construct valid analytic confidence intervals in the later event study regressions.

This matching approach leverages the richness of our data – the fact that we have the near-universe of businesses in the U.S. – to flexibly control for differences in characteristics at baseline that might drive differential trends. Appendix section C describes the matching algorithm in more detail, and Table A.2 provides a comparison of treated and control firms in the matched sample (columns 3 and 4) to treated and control firms the full sample (columns 1 and 2).[18] 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 in the BG sample that must restrict to firms that post ads in BG. The distribution of ads in BG are not targeted, but the matched control group does better on some of these, for instance, education requirements and the occupation distribution, compared to the full sample.

The full sample control group and the matched control group each have conceptual advantages and disadvantages. Differences across grant and non-grant firms leading up to the grant application are interesting in their own right. The full sample control comparison helps us better understand the nature of the skills gap problem and the ways in which the training firms have attempted to solve it, prior to training. Our goal is to describe these differences, as well as those post-training, compared to a reasonable control: firms of a similar age, size, and growth operating in similar markets and time periods. In contrast, the matched control helps us rule out alternative stories in which the post-grant firm outcomes are driven by differences in the types of firms which apply for training grants. For instance, if patterns reflect that firms tend to apply for training at a certain phase of their life-cycle, those should be picked up by our match on size trends. The matched sample provides the cleanest estimate of the added impact of these training grants apart from selection effects.

[18] As detailed in the appendix, we exclude from the matching analysis treated firms that do not achieve a sufficiently close match among the control firms – 12% of treated firms.

5.2 Results

Quantity of Employment and Vacancies

We begin with log employment as measured in the QCEW. Figure 7 plots event study coefficients and 90% confidence intervals for the full sample and for the matched sample. Appendix Table A.5 reports the coefficients and standard errors for these specifications, as well as specifications that add sector-year fixed effects. In the full sample specification, we see that grant-receiving firms are on a different growth trajectory prior to grant receipt. By construction, this gap closes when comparing to the matched sample.[19] Focusing on the matched sample, we see that firms grow steadily after receiving training. By 5 years after training receipt, firms are about 17% larger relative to the baseline matched control employment level of 143 workers.

Figure 7: Firms Grow After Training: Employment Event Studies

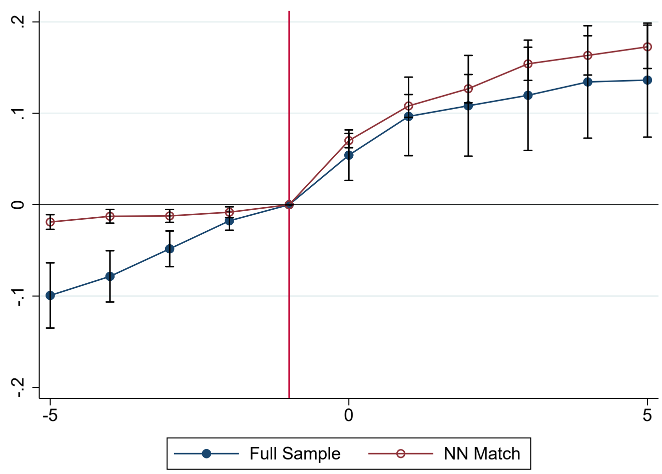

Notes: This figure reports coefficients estimated using equation 3, our event study regression, for log employment separately for the full sample and the matched sample. We control for establishment fixed effects, year fixed effects, and event dummy main effects. We plot coefficients on event dummies interacted with treatment. We also report the 90% confidence interval based on standard errors clustered at the state level.

We next examine the presence and quantity of vacancy postings to better understand whether firms had a stated preference for this growth (as opposed to passive hiring or changes in their separation rates). Panel A of Figure 8 and Appendix Table A.6 report the results of regressing an indicator for any postings in BG on the event study specification, which, in addition to helping to understand

[19] The coefficients in periods t-2 to t-5 are significantly different than the coefficient normalized to zero in period t-1, but they are not significantly different from each other and they are tiny in magnitude, relative to the impacts in the post-period.

any differential selection into the BG sample, can be viewed as extensive margin effects. As with log employment, the full sample exhibits growth in the likelihood of postings a vacancy in the pre-period, but, by construction, we have parallel trends in the matched sample. We see a sharp increase in any posting in the matched sample with grant recipients increasing their likelihood of online hiring by 6-7 pp off a baseline of 45% in the matched control.

We then next look at the intensive margin by regressing number of posts in BG on our event study specification. For this sample, we restrict to firms that match in that year to the BG sample, meaning they have posted at least one job ad online that BG was able to capture.[20] In Panel B of Figure 8 and Appendix Table A.7, we see similar patterns to the extensive margin; in the matched sample, grant firms post on average 19 more ads in BG five years post-grant than non-grant firms – conditional on posting at all – which is a 67% increase relative to the baseline.

Figure 8: Firms Increase Online Hiring After Training: Vacancy Event Studies

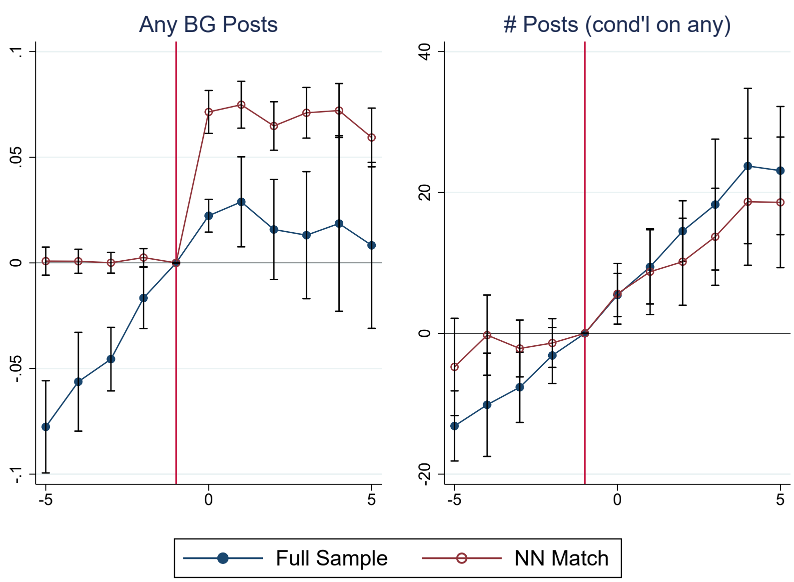

Notes: This figure reports coefficients estimated using equation 3, our event study regression, for an indicator for posting in BG in the year and for the number of posts in BG, conditional on having any in the year. We control for establishment fixed effects, year fixed effects, and event dummy main effects. We plot coefficients on event dummies interacted with treatment. We also report the 90% confidence interval based on standard errors clustered at the state level.

Furthermore, these event studies show that the change occurs gradually, reaching its highest point 5 years out. Almost all grants last for two years or less, with most being completed within a year

[20]For comparability, appendix figure A.6 reproduces QCEW results for log employment restricting to this same sample and finds similar magnitudes and patterns.

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 Employment and Vacancies

As vacancies and employees increase, we might think the characteristics of the jobs firms are hiring in has also changed. One possible outcome of the grants is that firms no longer need to include skill requirements in their job vacancies. Conversely, once the firm has built a workforce in the desired skill, it might need to hire for tasks that complement the trained workers leading to either an increase or a decrease in skill requirements depending on the type of complementarity. To get a better sense of employment composition, we exploit the rich detail in the BG vacancy data, using job ads as a proxy for how the workforce might be changing.

We start by looking at whether firms change which occupations they are hiring in. We examine the proportion of ads across the four occupation groups described above: Professional, Administrative, Low-Skill Service, and Production/Blue Collar. We estimate equation 3 at the establishment-year level for both the full sample and the nearest neighbor sample. To better understand how the distribution of vacancies has changed over time, we weight observations by the number of ads posted. Thus, results can be interpreted as impacts on the average vacancy of a treated firm, rather than the impact on the average treated firm. Figure 9 and Appendix Table A.8 reports the results of these regressions.

We find that when firms receive grants, the composition of their ads shifts away from professional occupations. Effects are statistically significant and the event study (top left panel of figure 9) shows that effects only appear both after grant receipt and after the training has occurred in both the full and the nearest neighbor specification. Five years post-grant receipt, grant recipients' hiring requests are 2 pp less likely to be in professional occupations relative to a baseline of 57 in the control firms. Since this decline does not occur until 3-4 years post receipt, it seems unlikely that effects are mechanically driven by training. Instead, firms are more likely to post in administrative white collar positions such as sales or office support staff (1.2 pp off a base rate of 19 percent). We see positive, though statistically insignificant, point estimates for the other two groups, suggesting that the shift away from professional occupations is distributed across all the other occupation groups.

We next test whether the ads ask for a higher or lower level of skill as measured by required

Figure 9: Ad Shares across Broad Occupation Groups (BG)

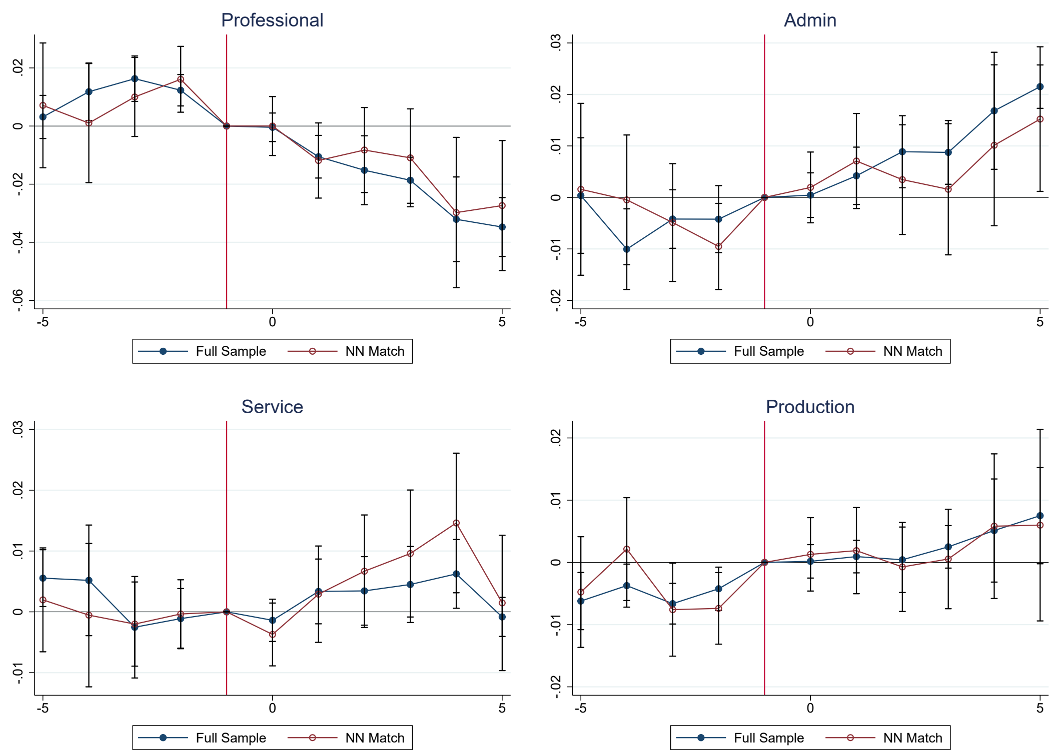

Notes: See figure 7 for regression specification information. Here we restrict to establishments that had BG postings in the year. Outcomes include the fraction of postings in each of four mutually exclusive and exhaustive occupation groups: professional, routine white collar, Low-skill service, and blue collar (see footnote 12).

education and experience. All regressions are again weighted by number of ads posted by the firm in a year. Focusing on the event studies for the preferred specification in Figure 10, several interesting patterns emerge.[21] First, we see that the composition of education requirements changes with training employers being 4.6 pp (about 10%) less likely to require a college degree, and these effects are statistically significant with no differential trend in the pre-period. While the level of requirement changes, the evidence for changes in any skill requirement is more mixed. Though the full sample regressions suggest that training firms require less experience and education directly post grant-receipt, this pattern disappears in the nearest neighbor sample where we match on pre-period employment trends. Together, these two figures imply that employers reduce education requirements – replacing a college requirement with a high school diploma – but do not eliminate skill requirements all together.

Figure 10: Skill Requirements (BG) Event Studies

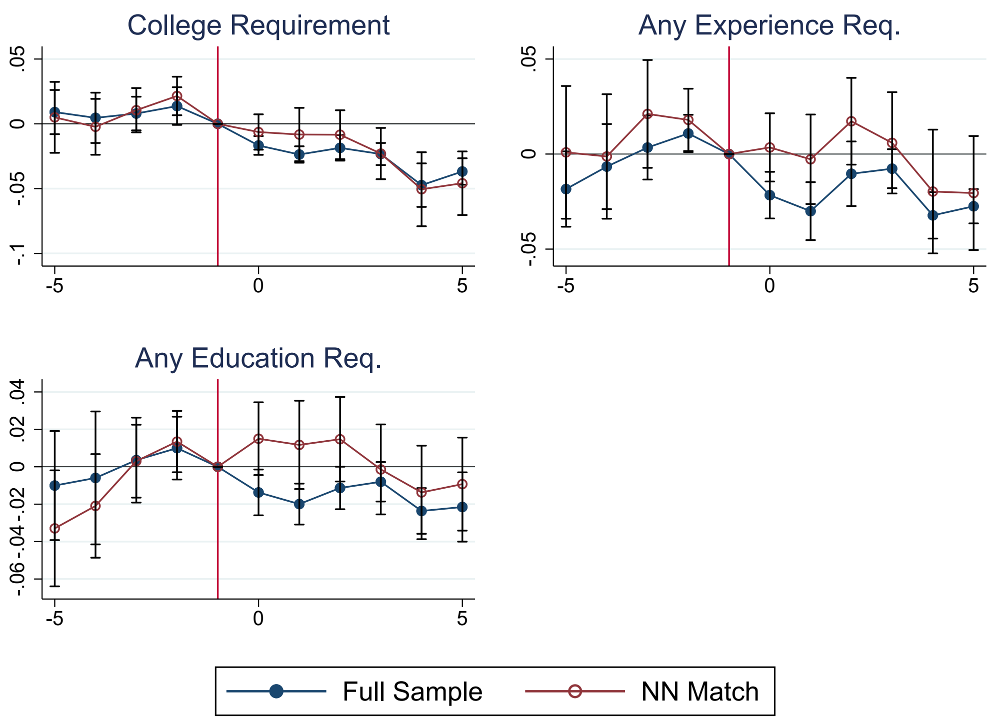

Notes: See figure 7 for regression specification information. Outcomes are the proportion of ads specifying the indicated skill requirement: any education, a college degree, experience.

The reduction in any education and experience requirements in the full sample show up almost immediately after grant receipt – suggesting they could be in part mechanical. A firm may anticipate

[21] Appendix Table A.8 reports the coefficients and standard errors for these specifications.

that it will train the newly hired workers and decide it does not need to impose skill requirements. However, when we look at the college requirement, this effect does not appear in both samples until three years post-training, which would reflect other changes in firms' recruiting strategies spurred by training rather than a mechanical effect. Some of the reduction in skill requirements could be driven by the fact that the firm has downskilled its occupation distribution as well. However, we have explored these results controlling for the distribution of occupations across the four aggregate categories and find similar results: the decreases in skill requirements remain.

Lastly, we look at the impacts of training on wages. A priori, the direction of this shift is not necessarily clear. On one hand, the training grants typically explicitly require firms to increase the wages they pay to trained employees. Training in general skills would also increase a worker's outside option, putting upward pressure on wages. On the other hand, once the firm has trained a subset of their workforce, they may then expand hiring in complementary positions. Given that we see a decline in hiring in higher paid professional positions and an increase in lower paid sales/administrative positions, we might expect this effect to decrease total wages per worker.

Figure 11: Wages per Worker are Unchanged Post-Training: Log Wage Event Studies

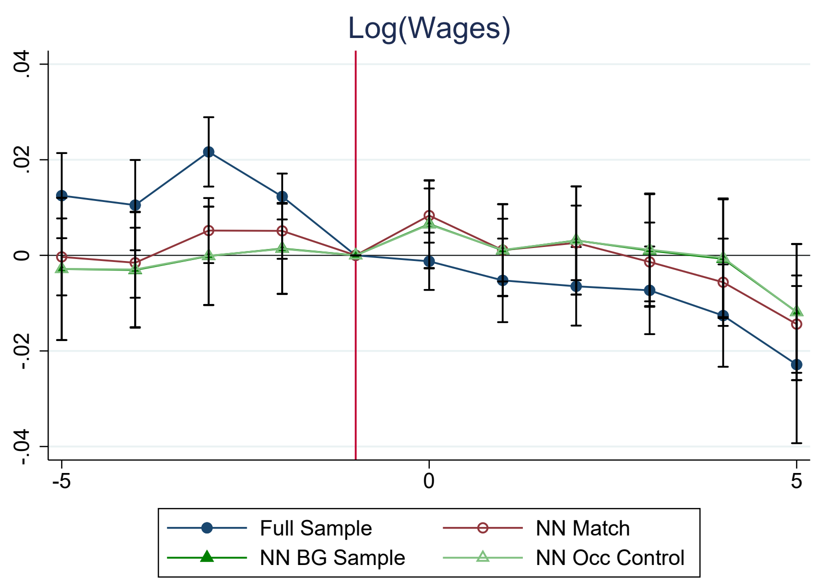

Notes: This figure reports coefficients estimated using equation 3, our event study regression, for the log of total wage bill in the year, controlling for log employment. See 7 notes for controls and standard error specification. We run these regressions separately for the full sample (navy), matched sample (red), and the matched sample restricted to firms with at least one post in BG (green). In the BG postings sample, we run a specification without (solid triangle) and with (hollow triangle) controls for proportion of ads in professional, sales, service, and production occupations.

Figure 11 and Appendix Table A.9 shows regression results for the log of the total wage bill for the firm. This regression also controls for log employment in the year to create an estimate of wage per worker.[22] We first look at wages for the nearest neighbor matched sample and see wages are stable in the pre-period and flat for the post period with a slight downward trend five years post-grant. Because changing skill composition of the firm may obfuscate a positive wage effect if the firm's hiring is primarily lower skill workers, we next want to test whether these results hold controlling for occupational mix. We can approximate changes in firm occupation mix by controlling for the proportion of ads hiring in each of the four previously described occupation groups (i.e., Professional, Sales/Admin, Service, and Production). We therefore also show regression results for the nearest neighbor match conditional on there being at least one post in BG without controls for occupational mix and with controls for occupational mix. The BG sample has slightly higher, but noisier coefficients, but we still cannot reject that there is a null effect of grant receipt on wages. The coefficients for the regressions with and without occupation controls overlap, suggesting that changes in skill composition is not enough to explain this null effect on wages.

5.3 Discussion

In this section, we have shown that, post-training, establishments grow both in terms of the number of employees and the number of vacancies. The composition of jobs also appears to have changed with training firms shifting away from professional occupations and away from explicit degree 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.

We can think of a few alternative hypotheses for why these changes occur. First, though we have done our best to find reasonable control groups, training is non-randomly assigned. We cannot rule out that training firms would have seen these outcomes even absent training. That does seem unlikely, however, given the consistency of results and the lack of pre-trends for outcomes in the matched sample. Clearly, there is something changing for training firms around the training period.

Second, training could have a real impact on production. Firms may have had a bottleneck in the production process and, once resolved, the firm is able to produce at scale and grow. This channel is consistent with the fact that a large fraction of hiring growth is in sales and administrative

[22] This estimate of log wage per worker is quite noisy. Wages 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. Thus firms experiencing heavy churn will appear to pay higher wages per worker. Also, we cannot distinguish wages of incumbents versus new hires whose wages might be more flexible.

positions. As production needs are resolved, the firm will wish to grow in tasks that are complementary to training, such as front-line sales positions – consistent with our results on occupational outcomes. Why do we see skill requirements decline? It could be that even within broad occupation categories, the tasks that complement training do not require as much skill. Alternatively, firms may have realized that training is a viable option for upskilling its workforce. They may back off of requirements they thought they needed for a wide range of positions, in favor of producing those skills in house.

The result is that 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).

6 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 help with employment prospects. Further, employers may be reluctant to pay to train workers themselves when they risk their investments being poached away.

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 have greater poaching risk, as measured by the concentration of vacancy postings. Finally, we find that grant participation facilitates growth, as measured by an increase in vacancies. This growth is disproportionately concentrated in lower skilled front-line positions. Even conditional on the changing composition of jobs, firms relax skill requirements in job postings.

Overall, our findings are inconsistent with place-based development motivations. In particular, we 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 mega-firms that might hold out-sized political power. 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. Firms are reluctant to train due to increased poaching risk and workers may lack the resources or awareness to find training on their own. 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 the firm who now hires in complementary jobs.

The stated goals of many of these state programs is to upskill the state's workforce and assist firms in weather out-of-state competition. We view our results as very positive from the perspective of the social planner and suggest these programs could be scaled up.

References

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.

Angell, M., S. Gold, J. S. Hastings, M. Howison, S. Jensen, N. Keleher, D. Molitor, and A. Roberts (2021). Estimating value-added returns to labor training programs with causal machine learning.

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.

Bartik, T. J. (2017). A new panel database on business incentives for economic development offered by state and local governments in the united states.

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.

Brunello, G. and M. De Paola (2008). Training and economic density: Some evidence form italian provinces. *Labour Economics 15*(1), 118–140.

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].

Commonwealth Corporation (2024). Workforce training fund program: General program guidelines. `https://commcorp.org/subprogram/wtfp-general-program-guidelines/` [Accessed: September 10, 2024].

Dalton, M. R., L. B. Kahn, and A. I. Mueller (2023). Do online job postings capture job vacancies? *Mimeo*.

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.

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.

Goldin, C. and L. Katz (2008). *The Race betwee Education and Technology*. Harvard University Press.

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.

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.

Hollenbeck, K. (2008). Is there a role for public support of incumbent worker on-the-job training? Working Paper 08-138, Upjohn Institute.

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.

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, 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.

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.

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.

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. (2020). Bidding for firms: Subsidy competition in the u.s. Technical report, Columbia GSB Working paper.

Stevens, M. (1996). Transferable training and poaching externalities.

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.

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.

Appendix

A 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, standarizing 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. (2023) to bring in ad characteristics.[23] 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.

Table A.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. Earlier work 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).

[23] Note, 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.

B Categorization of Training Grant Descriptions

For a sub-set of the states in our sample, 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 A.7 shows the system-level prompt that we supplied to GPT for the classification task, and Figure A.8 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. (2023), 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 submitted to California, Kentucky, Massachusetts , New Hampshire, and New Jersey. 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 for 64 training plans.

To give a concrete example, the firm depicted in Figure A.8 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.

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. The "reference year" in this matching process is the year of first grant receipt for treated firms and the randomly assigned placebo year (see section 3.3) 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 log employment in periods $t-1$ to $t-5$ relative to the reference year and indicators for having any job-posting activity in Burning Glass in $t-1$ to $t-5$. For firms with missing log employment for some years of the pre-period we fill in a value of -1000, 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. We use the BallTree nearest neighbor matching algorithm, implemented in SciKitLearn, to match efficiently. 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 in log employment over the pre-period (including any -1000 missing indicators) is greater than 0.16. 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.

D Appendix Figures and Tables

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

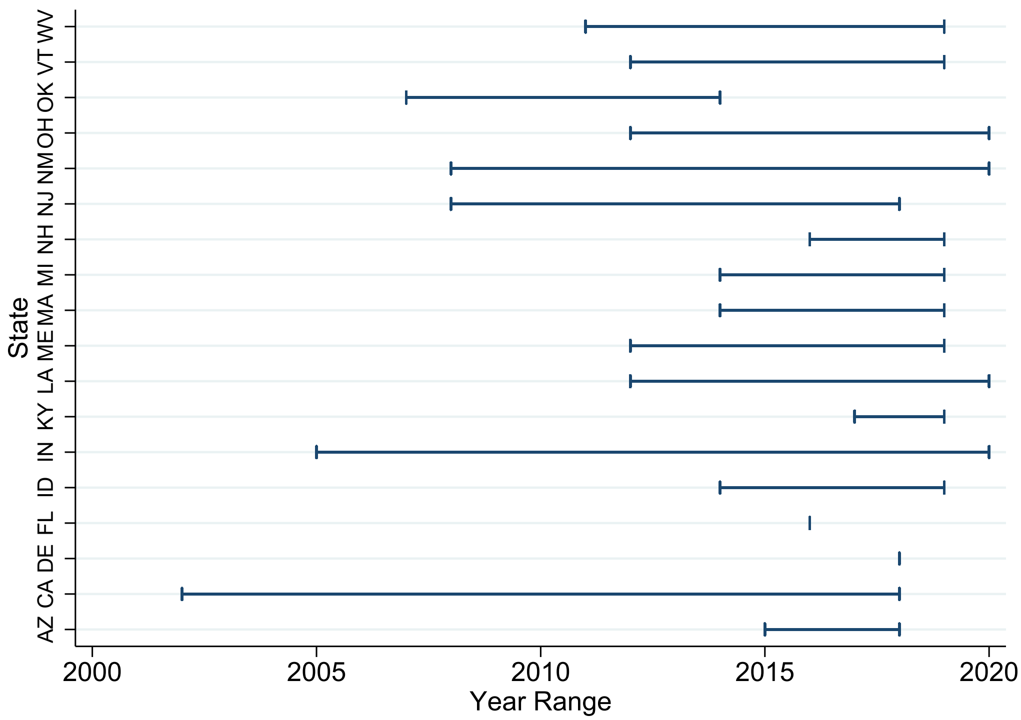

Notes: We summarize the range of years for which grant data are available for a given state. Grant data are assembled by the authors by reviewing state department of labor websites for training programs characterized by public funds flowing to individual firms to train their own workers. We include data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

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

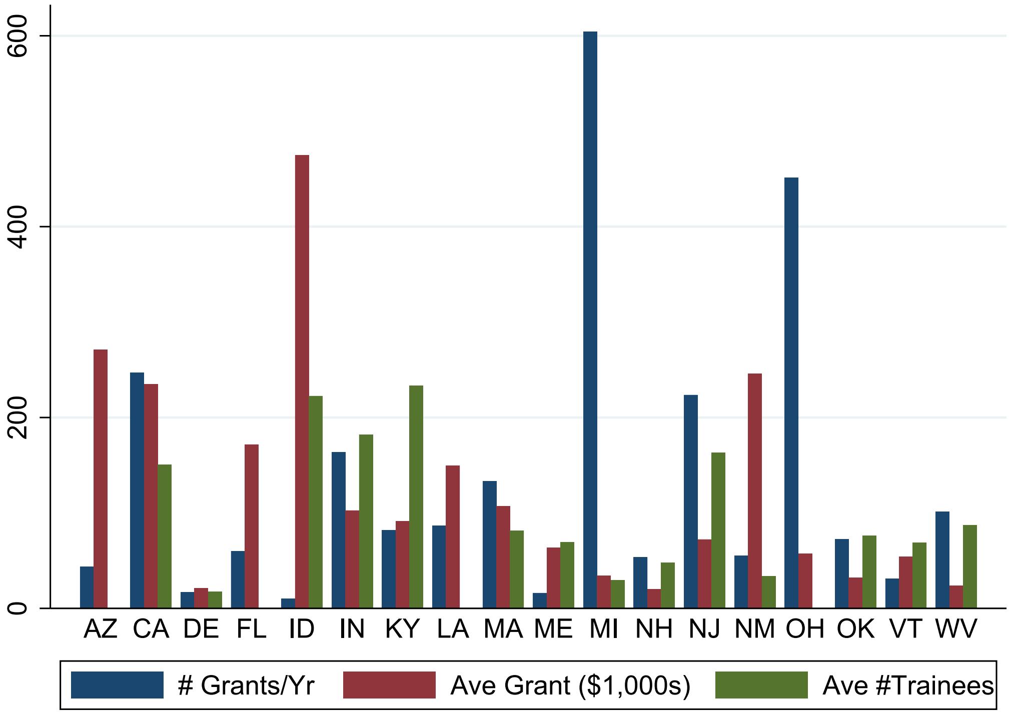

Notes: We plot characteristics of grants by state for the matched sample of grants (see table A.1). We report unweighted means of the number of grants per year, grant dollars and number of trainees. The latter is unavailable in a small number of states.

Figure A.3: Training Grants and Market Concentration by State-Level Competitiveness

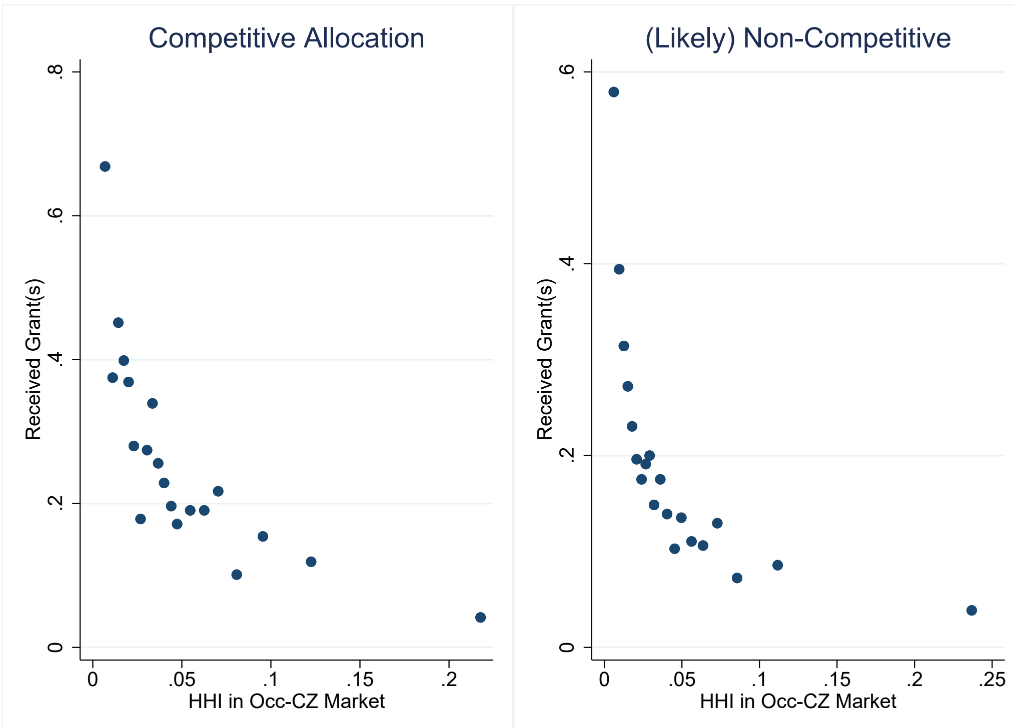

Notes: See figure 5. Here we split states into those with a rigorous scoring rubric and competitive selection process (left) and those with no apparent rigor (right). The former include Arizona, Kentucky, Louisiana, Michigan and New Hampshire.

Figure A.4: Training Grants and Market Concentration: Robustness to Number and Size of Grants

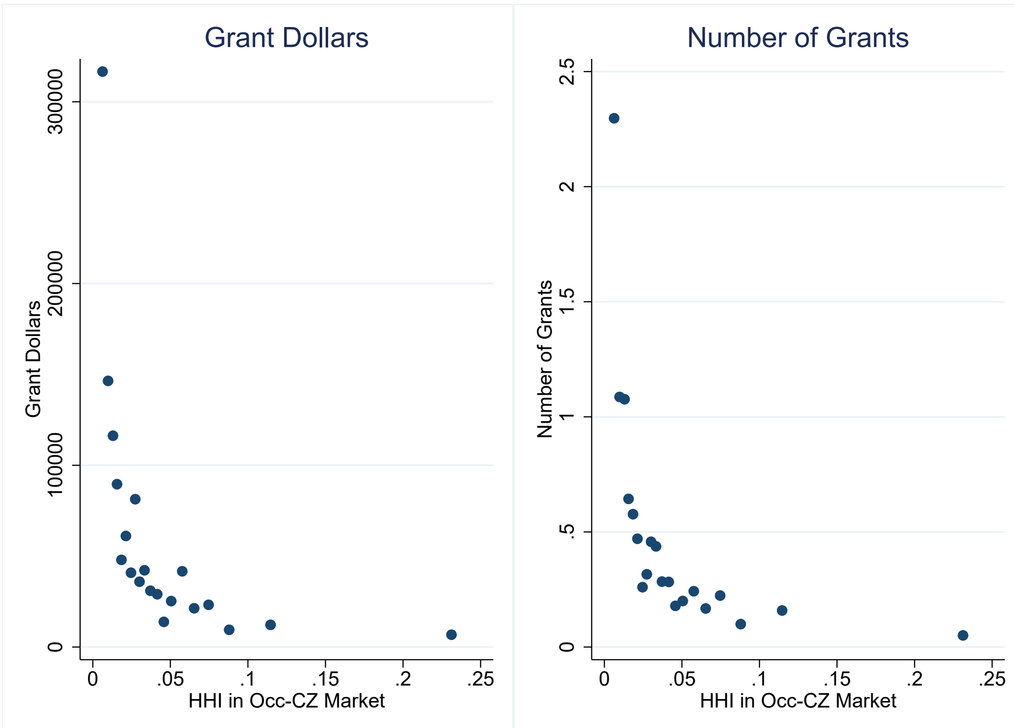

Notes: See figure 5. Here we plot bin scatters of total grant dollars or total number of grants (both including zeros) in a market on the concentration of vacancies. Markets are defined at the CZ-by-three-digit occupation level and concentration is the HHI of job vacancies posted in the market

Figure A.5: Training Grants and Market Concentration using industry-by-CZ Market Definition

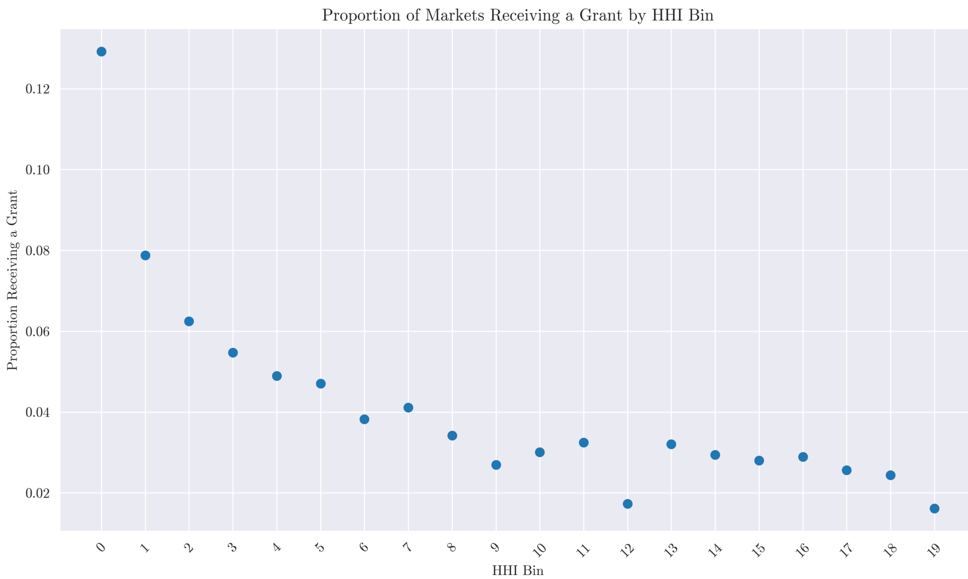

Notes: See figure 5. We divide markets (CZ-by-two-digit NAICS industry pairs) into 20 equally-sized bins based on the HHI of job vacancies posted in the market (see equation 2). We then plot the share of markets that received any grants within each bin.

Figure A.6: QCEW Outcomes Restricted to BG-Matched Sample

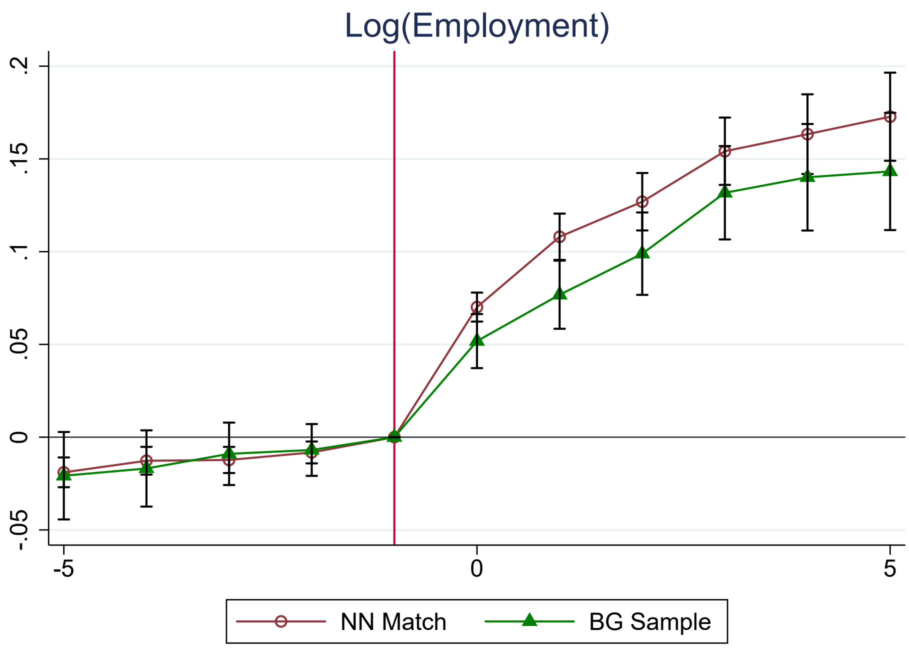

Notes: See figure 7 for regression specification details. The samples in this figure restrict to establishments that post at least one ad collected by BG over our sample period. We also report the 90% confidence interval based on standard errors clustered at the state level.

Figure A.7: 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.

Figure A.8: Example of Training Description

OVERVIEW

For over 60 years, Arrow Sign Company (Arrow) has manufactured electric and architectural signs. The Company's products are used primarily for brand identification and business location visibility. Arrow's 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, Arrow provides services ranging from initial design concepts to detailed plans. The Company also provides fabrication, installation and maintenance of their products. Arrow customers include major hotels, property management companies, building owners, shopping centers, and general contractors.

Arrow uses Computer-Aided Design and Computer-Aided Manufacturing software integrated for design and fabrication. This technology provides Arrow with a competitive advantage. In order to meet growing customer demands and preserve its market share, Arrow 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 Arrow Sign 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.

Table A.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	90837	115976	87633	108667
	(191109)	(188184)	(237387)	(167363)	(276163)
	N=13249	N=12564	N=685	N=10650	N=1914
# Trainees	94.39	92.67	122.89	96.34	74.72
	(210.91)	(206.33)	(274.82)	(214.73)	(157.67)
	N= 9966	N=9400	N= 566	N=7808	N=1592
# New Hires	57.60	58.05	50.04	60.85	45.91
	(1443.53)	(1485.58)	(130.32)	(1645.58)	(182.68)
	N= 5335	N= 5035	N= 300	N=4092	N=943
Grant Dollars per Trainee	2240.0	2259.8	1904.4	2040.6	3365.4
	(4522.2)	(4600.7)	(2866.1)	(3676.8)	(7635.9)
	N= 9645	N= 9107	N= 538	N=7600	N=1507
Grant year	2015.2	2015.2	2015.3	2015.3	2014.8
	(2.6)	(2.6)	(2.6)	(2.5)	(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 procedure in Dalton et al. (2023).

Table A.2: Summary Statistics for Grant and Non-Grant Firms

	All Firms		NN Matched	
	Grant	Non-Grant	Grant	Non-Grant
Panel A:	QCEW Sample			
Employment	209.76	24.66	148.85	143.27
Wage per Worker	62093	43153	61516	57575
Age	17.02	14.01	17.02	17.18
Annual Growth Rate	0.075	0.053	0.071	0.048
Ever Posted in BG	0.82	0.38	0.80	0.72
Posted in t-1	0.49	0.13	0.46	0.45
# Establishments	8495	1514273	7505	7505
Panel B:	BG Matched Sample			
Establishment-Level Characteristics:				
Employment	234.69	42.55	165.15	168.18
Wage per Worker	63294	49140	62723	60900
Age	17.36	14.41	17.36	17.95
Annual Growth Rate	0.078	0.061	0.072	0.044
# Posted in t-1	0.54	0.30	0.52	0.56
# BG Postings in t-1	41.05	15.67	33.93	28.74
# Establishments	6938	577743	6010	5425
Ad-Weighted Characteristics:				
Education Req.	0.71	0.55	0.70	0.65
College Req.	0.46	0.28	0.44	0.42
Experience Req.	0.60	0.47	0.59	0.55
Professional Occ.	0.65	0.46	0.63	0.57
Admin Occ.	0.16	0.25	0.17	0.19
Service Occ.	0.04	0.14	0.04	0.08
Production Occ.	0.11	0.12	0.11	0.12
Computer Req.	0.33	0.25	0.33	0.30
Cognitive Req.	0.41	0.29	0.40	0.38
# Ads Total	153,169	2,675,878	105,364	87,956

Notes: This table reports characteristics for grant and non-grant recipients. Panel A uses the entire QCEW matched sample, while panel B restrict so the BG sample. 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 A.3: Training firms and market characteristics: 2-digit Industry-by-CZ

Dependent Variable	Any Grants Received (mean = 0.195)			
	(1)	(2)	(3)	(4)
HHI	-1.182***	-1.174***	-1.130***	-1.166***
	(0.149)	(0.143)	(0.158)	(0.153)
HHI^2	1.194***	1.145***	1.121***	1.175***
	(0.216)	(0.197)	(0.231)	(0.217)
CZ unemp rate	-0.735	-0.973**	-0.610	-0.700
	(0.446)	(0.453)	(0.449)	(0.486)
New Market	-0.156**	-0.159**	-0.215***	-0.159*
	(0.073)	(0.069)	(0.063)	(0.077)
Employment (1,000s)	0.042***	0.041***	0.041***	0.041***
	(0.005)	(0.006)	(0.005)	(0.005)
Wage ($100s)	0.891***	0.986***	1.154***	0.919***
	(0.174)	(0.176)	(0.287)	(0.185)
Fraction in small firms <10	-0.122*	-0.142***	-0.104*	-0.105*
	(0.059)	(0.048)	(0.055)	(0.056)
Fraction in large firms >500	0.154**	0.219***	0.174**	0.166**
	(0.067)	(0.068)	(0.067)	(0.065)
Emp growth			-0.066	-0.075
			(0.048)	(0.050)
Wage growth			-0.008	-0.007
			(0.081)	(0.074)
Leave-out State Emp			-0.024*	
			(0.012)	
Leave-out State Wage			-0.005	
			(0.003)	
State Emp Growth			0.249***	
			(0.084)	
State Wage Growth			-0.028	
			(0.113)	
Leave-out Region Emp				-0.002
				(0.003)
Leave-out Region Wage				-0.000
				(0.000)
Region Emp Growth				0.317***
				(0.071)
Region Wage Growth				-0.053
				(0.108)
Observations	17,229	17,228	17,229	17,229
R-squared	0.297	0.342	0.301	0.300
Occ, State-by-Year FEs	X	X	X	X
Occ-by-year,Occ-by-State		X		

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

Notes: Observations are 2-digit industry-by-CZ-by-year. HHI, Employment, Wages, and fraction of employment in small and large firms are industry-by-CZ averages from 2010-12, measured in County Business Patterns. Emp and wage growth are the rate of change in 2012 from 2010 for the ind-by-CZ. The 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 9,856 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".

Table A.4: Robustness to alternative measures of market concentration

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Received Grant (Mean: .21)					
	Occupation mkt		Industry mkt			
HHI	-2.208***		-1.034***			
	(0.513)		(0.119)			
HHI^2	5.432***		1.056***			
	(1.439)		(0.176)			
Share of ads to top 3 Firms		-0.555***		-0.423***		
		(0.141)		(0.030)		
Share of Emp in top 3 Firms					-0.283***	
					(0.081)	
Industry Tightness						0.304***
						(0.050)
Employment (1,000s)	0.076***	0.072***	0.042***	0.041***	0.044***	0.048***
	(0.005)	(0.005)	(0.007)	(0.006)	(0.008)	(0.008)
Wage ($100s)	0.733**	0.679**	1.691***	1.538***	1.843***	1.680***
	(0.285)	(0.298)	(0.339)	(0.329)	(0.308)	(0.355)
Fraction in small firms <10			-0.187*	-0.162*	-0.138	-0.189*
			(0.090)	(0.092)	(0.097)	(0.099)
Fraction in large firms >500			0.217**	0.225**	0.304***	0.201**
			(0.084)	(0.083)	(0.082)	(0.094)
Observations	13,860	13,860	9,852	9,852	9,852	9,756
R-squared	0.399	0.400	0.355	0.360	0.348	0.352
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 2 and A.3. Regression observations restricted to 2013-2019 and restrict to markets with at least 50 ads from 2010-12. Industry tightness is the number of jobs posted in BG divided by 100 times the number of unemployment workers who previously worked in the industry as measured in the ACS.

Table A.5: Event Study Coefficients: Log Employment (QCEW)

Event Time	Full Sample	Full Sample	NN Sample	NN Sample
t-5	-0.099	-0.124	-0.019	-0.019
	(0.022)	(0.012)	(0.005)	(0.005)
t-4	-0.078	-0.095	-0.013	-0.013
	(0.017)	(0.011)	(0.005)	(0.005)
t-3	-0.048	-0.060	-0.012	-0.012
	(0.012)	(0.009)	(0.004)	(0.004)
t-2	-0.018	-0.023	-0.008	-0.008
	(0.006)	(0.006)	(0.004)	(0.004)
t-1	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
t=0	0.054	0.056	0.070	0.070
	(0.017)	(0.016)	(0.005)	(0.005)
t+1	0.097	0.099	0.108	0.108
	(0.026)	(0.026)	(0.008)	(0.008)
t+2	0.108	0.112	0.127	0.127
	(0.033)	(0.033)	(0.009)	(0.009)
t+3	0.120	0.126	0.154	0.154
	(0.037)	(0.036)	(0.011)	(0.011)
t+4	0.134	0.126	0.163	0.163
	(0.037)	(0.038)	(0.013)	(0.013)
t+5	0.136	0.138	0.173	0.172
	(0.038)	(0.037)	(0.014)	(0.014)
Firm FE	X	X	X	X
Year FE	X		X	
Sector-State-Year FE		X		X
R-squared	0.8887	0.8897	0.9235	0.9246
Observations	14447856	14447856	147958	147958

Notes: This table reports event study coefficients for regression specification 3 with log employment as the outcome. Column 1 and 2 use the full-sample control; Column 3 and 4 use the nearest neighbor matched control. Odd columns correspond to event studies graphed in Figure 7 and even columns add two-digit industry by state by year FE.

Table A.6: Event Study Coefficients: Any BG Posts

Event Time	Full Sample	Full Sample	NN Sample	NN Sample
t-5	-0.078	-0.074	0.001	0.001
	(0.013)	(0.012)	(0.004)	(0.004)
t-4	-0.056	-0.059	0.001	0.001
	(0.014)	(0.013)	(0.003)	(0.003)
t-3	-0.046	-0.046	0.000	0.000
	(0.009)	(0.009)	(0.003)	(0.003)
t-2	-0.017	-0.017	0.003	0.003
	(0.009)	(0.009)	(0.003)	(0.003)
t-1	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
t=0	0.022	0.018	0.071	0.071
	(0.005)	(0.003)	(0.006)	(0.006)
t+1	0.029	0.021	0.075	0.075
	(0.013)	(0.012)	(0.007)	(0.007)
t+2	0.016	0.020	0.065	0.065
	(0.014)	(0.015)	(0.007)	(0.007)
t+3	0.013	0.018	0.071	0.072
	(0.018)	(0.018)	(0.007)	(0.007)
t+4	0.019	0.022	0.072	0.072
	(0.025)	(0.025)	(0.008)	(0.008)
t+5	0.008	0.011	0.059	0.060
	(0.024)	(0.023)	(0.008)	(0.008)
Firm FE	X	X	X	X
Year FE	X		X	
Sector-State-Year FE		X		X
R-Squared	0.5258	0.5306	0.5612	0.5670
Observations	13213284	13213284	131628	131628

Notes: This table reports event study coefficients for regression specification 3 with an indicator for having at least one post in BG as the outcome. Column 1 and 2 use the full-sample control; Column 3 and 4 use the nearest neighbor matched control. Odd columns correspond to event studies graphed in Figure 8 Panel A and even columns add two-digit industry by state by year FE.

Table A.7: Event Study Coefficients: Number of Vacancies (Cond. on Any)

Event Time	Full Sample	Full Sample	NN Sample	NN Sample
t-5	-13.154	-16.148	-4.771	-4.747
	(3.026)	(3.253)	(4.201)	(4.204)
t-4	-10.150	-12.884	-0.255	-0.366
	(4.458)	(4.538)	(3.462)	(3.481)
t-3	-7.659	-8.951	-2.153	-2.242
	(3.035)	(3.245)	(2.454)	(2.468)
t-2	-3.149	-3.718	-1.378	-1.384
	(2.415)	(2.533)	(2.098)	(2.107)
t-1	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
t=0	5.441	5.450	5.619	5.685
	(1.867)	(1.823)	(2.620)	(2.630)
t+1	9.435	9.411	8.745	8.926
	(3.203)	(3.304)	(3.702)	(3.724)
t+2	14.513	15.081	10.169	10.385
	(2.618)	(2.732)	(3.756)	(3.770)
t+3	18.284	19.078	13.713	13.796
	(5.651)	(5.620)	(4.188)	(4.155)
t+4	23.757	24.349	18.685	18.819
	(6.705)	(6.647)	(5.480)	(5.489)
t+5	23.105	24.565	18.600	19.128
	(5.533)	(5.634)	(5.638)	(5.646)
Firm FE	X	X	X	X
Year FE	X		X	
Sector-State-Year FE		X		X
R-Squared	0.6787	0.6795	0.6441	0.6482
Observations	1982990	1982990	59810	59810

Notes: This table reports event study coefficients for regression specification 3 with number of postings in BG as the outcome. Column 1 and 2 use the full-sample control; Column 3 and 4 use the nearest neighbor matched control. Odd columns correspond to event studies graphed in Figure 8 Panel B and even columns add two-digit industry by state by year FE.

Table A.8: Event Study Coefficients: Prop. in Skill Group

Event Time	Professional	Sales	Service	Production	College	Any Ed	Any Exp
t-5	0.007	0.002	0.002	-0.005	0.005	-0.033	0.001
	(0.013)	(0.010)	(0.005)	(0.005)	(0.017)	(0.019)	(0.021)
t-4	0.001	0.000	-0.001	0.002	-0.002	-0.021	-0.001
	(0.012)	(0.008)	(0.007)	(0.005)	(0.013)	(0.017)	(0.020)
t-3	0.010	-0.005	-0.002	-0.008	0.011	0.003	0.021
	(0.008)	(0.007)	(0.004)	(0.005)	(0.010)	(0.012)	(0.017)
t-2	0.016	-0.010	0.000	-0.007	0.022	0.013	0.018
	(0.007)	(0.005)	(0.003)	(0.003)	(0.009)	(0.010)	(0.010)
t-1	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
t=0	0.000	0.002	-0.004	0.001	-0.006	0.015	0.003
	(0.006)	(0.004)	(0.003)	(0.004)	(0.008)	(0.012)	(0.011)
t+1	-0.012	0.007	0.003	0.002	-0.008	0.012	-0.003
	(0.008)	(0.006)	(0.005)	(0.004)	(0.013)	(0.014)	(0.014)
t+2	-0.008	0.003	0.007	-0.001	-0.008	0.015	0.017
	(0.009)	(0.006)	(0.006)	(0.004)	(0.011)	(0.014)	(0.014)
t+3	-0.011	0.002	0.010	0.001	-0.023	-0.001	0.006
	(0.010)	(0.008)	(0.006)	(0.005)	(0.012)	(0.015)	(0.016)
t+4	-0.030	0.010	0.015	0.006	-0.050	-0.014	-0.020
	(0.016)	(0.010)	(0.007)	(0.007)	(0.017)	(0.015)	(0.020)
t+5	-0.027	0.015	0.001	0.006	-0.046	-0.009	-0.021
	(0.014)	(0.009)	(0.007)	(0.009)	(0.015)	(0.015)	(0.018)
Firm FE	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X
Observations	59810	59810	59810	59810	59810	59810	59810

Notes: This table reports event study coefficients for regression specification 3 for proportion of ads asking for skill types including Professional, Sales, Service, and Production occupations, college degree requirement, any education requirement, and any experience requirement. All are in the nearest neighbor matched sample and include firm and year FE. These correspond to the nearest neighbor event studies in Figures9 and 10.

Table A.9: Event Study Coefficients: Log Wages

Event Time	NN Match	NN BG Match	NN BG Match
t-5	0.000	-0.003	-0.003
	(0.005)	(0.009)	(0.009)
t-4	-0.002	-0.003	-0.003
	(0.004)	(0.007)	(0.007)
t-3	0.005	0.000	0.000
	(0.004)	(0.006)	(0.006)
t-2	0.005	0.001	0.001
	(0.004)	(0.006)	(0.006)
t-1	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
t=0	0.008	0.007	0.006
	(0.003)	(0.006)	(0.006)
t+1	0.001	0.001	0.001
	(0.004)	(0.006)	(0.006)
t+2	0.003	0.003	0.003
	(0.005)	(0.007)	(0.007)
t+3	-0.001	0.001	0.001
	(0.005)	(0.007)	(0.007)
t+4	-0.006	-0.001	-0.001
	(0.006)	(0.008)	(0.008)
t+5	-0.014	-0.012	-0.012
	(0.006)	(0.009)	(0.009)
Firm FE	X	X	X
Year FE	X	X	X
Occ Control			X
R Squared	0.9864	0.9903	0.9903
Observations	147953	59810	59809

Notes: This table reports event study coefficients for regression equationb3 for log wages. Column 1 is the full nearest neighbor matched sample; column 2 is restricted to the matched pairs conditional on appearing in BG; and column 3 adds controls for the proportion of ads in Profesional, Sales, Service, and Production sectors. These correspond to the event studies in Figure 11.