INTERNATIONAL MONETARY FUND

The Dynamic Effects of Local Labor Market Shocks on Small Firms in The United States

Philip Barrett, Sophia Chen, Li Lin, Anke Weber

WP/24/63

2024
MAR

WORKING PAPER

 WP/24/63

IMF Working Paper
Western Hemisphere Department

The Dynamic Effects of Local Labor Market Shocks on Small Firms in The United States
Prepared by Philip Barrett, Sophia Chen, Li Lin, and Anke Weber*

Authorized for distribution by Nigel Chalk
March 2024

ABSTRACT: We use payroll data on over 1 million workers at 80,000 small firms to construct county-month measures of employment, hours, and wages that correct for dynamic changes in sample composition in response to business cycle fluctuations. We use this to estimate the response of small firms' employment, hours and wages following tighter local labor market conditions. We find that employment and hours per worker fall and wages rise. This is consistent with the predictions of the response to a demand shock in the well-known "jobs ladder" model of labor markets. To check this interpretation, we show our results hold when instrumenting for local demand using county-level Department of Defense contract spending. Correction for dynamic sample bias is important -- without it, the hours fall by only one third as much and wages increase by double.

JEL Classification Numbers:	D22, E24, E31, J2, J23, J63, J43
Keywords:	small firms; wages; hours, firm heterogeneity; private-sector establishment-level data; business cycle.

Author's E-Mail Address:	pbarrett@imf.org, ychen2@imf.org, llin@imf.org, aweber@imf.org

* We thank Nigel Chalk, Lydia Cox, Jason Greenberg, Yueling Huang, Do Lee, Ippei Shibata and Philippe Wingender for helpful suggestions and advice.

1 Introduction

The propagation of aggregate demand shocks to labor markets is a first order question in many macroeconomic policy issues. For example, labor markets play a key role in New Keynesian models of how monetary policy affects aggregate demand—at least in the short run (e.g., Christiano et al., 2016; Galí et al., 2011; Blanchard and Galí, 2010). As such, understanding the mechanisms by which demand for labor can affect the dynamics of employment, hours, and wages is an important topic for macroeconomists.

Early theories of how aggregate demand shocks transmitted via labor markets tended to treat labor as monolithic and supplied without frictions (e.g., Kydland and Prescott, 1982). Further work then integrated search-and-matching models of the labor market to the macro economy (Mortensen and Pissarides, 1994; Shimer, 2005; Pissarides, 2009; Mortensen, 2011). More recent macro-labor models have developed a "job ladder" theory, wherein the search for matches involves not only workers and firms but also encompasses a hierarchical structure of these matches. Larger, more productive firms produce a larger surplus available to be split with workers, raising their wages and making them more attractive to employees. This produces a "ladder" of jobs, with the best jobs (and the largest firms) at the top, motivating workers to strive for upward mobility. One of the most well-known expositions of this theory is Moscarini and Postel-Vinay (2016a, 2012).

One interesting prediction of the job-ladder model is that changing business conditions can have opposing effects on small firms. On the one hand, increased aggregate demand implies increased opportunities for small firms to expand, leading to higher levels of hiring and wages. On the other hand, the relatively low position of small firms on the job ladder may counteract these benefits. If workers vary by quality—and if there is already assortive matching—then larger, more productive firms will prefer to poach employees from their smaller competitors during expansions rather than hiring from the less-productive unemployed.[1] The lower a firm's position on the jobs ladder, the more exposed they are to worker poaching: there are more poachers above them and fewer poachees below. If the job ladder channel is strong enough, the smallest firms may experience negative overall effects with increased aggregate demand.

A simple calculation based on aggregate evidence hints at a non-trivial role for this mechanism. The job ladder theory predicts that cyclical variation in net relative to gross hiring should be smaller at smaller firms. This is because in expansions, increased competition from larger firms forces smaller firms to do more gross hiring to maintain a given level of staffing (replacing the staff lost to larger firms). In contractions, this phenomenon should reverse; the flood of higher-quality workers shed by higher-paid firms allows small firms to

[1] If workers are all of equal quality, larger firms will have no preference between hiring unemployed workers and poaching smaller firms' employees.

upgrade their workforce via large gross hiring and firing, even if net job changes are small. Figure 1 tests this relationship in aggregate data, showing that the standard deviation of net job growth as a proportion of gross hiring does indeed increase with firm size.[2]

Of course, there are other interpretations of these aggregate data. That is where this paper comes in. By using a rich dataset of millions of payroll records at small firms—which in 2022 accounted for 45 percent of aggregate private sector employment (CEA, 2023)—we aim to describe carefully the impact of business cycles on small firms' labor usage specifically. In doing so, we aim to test the mechanism outlined above and, by extension, evaluate (at least one part of) a leading theory of how labor markets work.

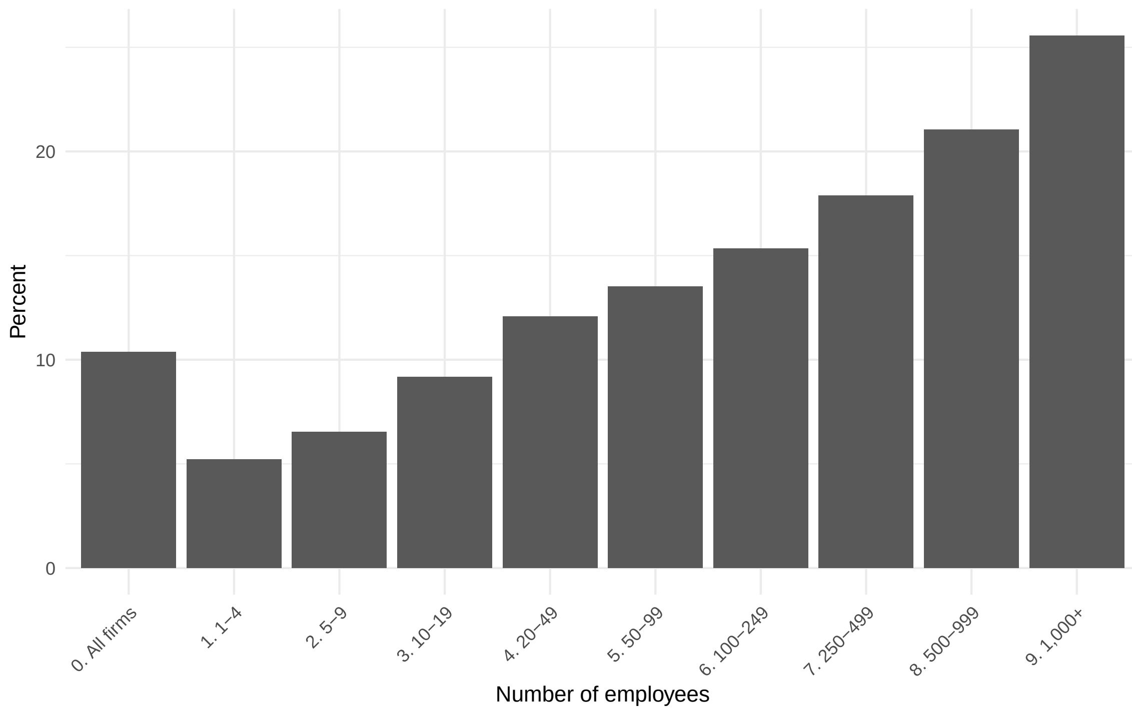

Figure 1: Net job growth standard deviation by firm size

Figure shows the standard deviation of net job growth $\sigma(g_{i,t})$ by firm size group i. Growth is computed as the ratio of net job growth to the sum of gross job gains and losses. Data on net and gross job flows come from the BLS business employment dynamics, quarterly during 1992Q1-2022Q3.

We use proprietary micro data from Homebase, a payroll provider for small businesses that provides broad-ranging information on hours and wages for nearly 9 million workers at 1 million firms. The Homebase data is particularly well-suited to answering the questions we pose because it almost uniquely covers small, low-wage firms. The median firm in our

[2]Data are taken from the U.S. Bureau of Labor Statistic's (BLS) Business Employment Dynamics. As is standard, net job growth is defined for size group i as $g_{i,t} = (gains_{i,t} - losses_{i,t})/(gains_{i,t} + losses_{i,t})$.

data has fewer than 7 employees, paying an average hourly wage of around $12, with almost exclusively casual or part-time employees.[3]

Our key methodological novelty is that we use the granularity of the Homebase data to construct county-month measures of employment, hours, and wages which correct precisely for dynamic composition bias. A key econometric challenge to estimating the dynamic responses of labor market variables to shocks is that the composition of aggregate data, even at the county-month level, may change systematically in response to shocks. For example, if incoming hires are intrinsically different from existing workers—as the jobs ladder theory suggests—impulse responses for wages or hours will not have a consistent interpretation across all horizons. By judicious choice of the sample of worker-firm observations at each horizon, we construct impulse responses which are robust to this form of dynamic composition bias.

Our main empirical findings are consistent with the jobs ladder theory. We construct a proxy indicator for county-month labor market tightness by taking the ratio of vacancies posted on the website Indeed.com to local unemployment. Tighter labor markets—which one would usually associate with strong aggregate demand—are tough for small firms: We find that an increase in local labor market tightness is followed by the firms in our dataset reducing employment and hours per worker, even though they increase wages. These effects persist for at least a year and are statistically significant. The number of firms and worker-firm matches observed also decreases, consistent with higher firm failure rates, although this could be due to other factors. Our composition correction methodology is also important for our findings. Without it, the impulse responses would paint a much rosier picture for small firms, suggesting that they barely reduce hours and can increase wages by almost twice as much. The composition correction also sheds some insight into the working conditions of those who move during periods of labor market tightness, with those leaving having lower hours and wages on average than those remaining, and those joining getting (unsurprisingly) higher wages.

We show that these results hold across firm size and industry, as well as before, during, and after the COVID pandemic. Finally, we investigate in some depth the interpretation of our results, arguing that they correspond to local demand shocks. To verify this, we use contract-level data from USASpending to construct country-month procurement spending by the Department of Defense procurement spending[4] and use this as an instrumental variable to isolate plausible local aggregate demand shocks. The results look much the same as the baseline.

[3]Although this dataset provides a detailed view into small firms' hiring and labor costs, one drawback is that we cannot trace employees across firms.

[4]https://www.usaspending.gov/

Literature. This paper is related to a number of important literature. An prominent line of research analyzes cross-sectional differences in the response of firms to aggregate shocks, paying close attention to firm size. Gertler and Gilchrist (1994) find that small firms are approximately twice as responsive to monetary shocks as large firms in terms of inventory demand and attribute this to financial frictions that can amplify the response of the economy to aggregate shocks. Using the U.S. Census Bureau's Quarterly Financial Report (QFR), a survey that collects income statements and balance sheets of manufacturing, retail, and wholesale trade firms, Crouzet and Mehrotra (2020) study the cyclical sensitivity of small firms' sales growth and find them to be more sensitive to aggregate fluctuations. Morales-Jiménez (2021) introduces workers' noisy information about the state of the economy and heterogeneous firms and demonstrates a pro-cyclical reallocation of employment from low- to high-paying firms, with high paying firms being larger and more productive. Most closely related to our results on employment is the work by Moscarini and Postel-Vinay (2012). The authors study the differential growth rate of employment between large and small U.S. firms using Census Bureau's Business Dynamic Statistics (BDS). Large employers are found to on net destroy proportionally more jobs relative to small employers when unemployment is above trend, late in and right after a typical recession, and create more when unemployment is below trend, late in a typical expansion. Moscarini and Postel-Vinay (2016a) formalize these insights through a model with a labor search mechanism. The key finding is that large firms are typically more productive, can pay more, and thus can successfully poach workers from smaller competitors. This makes their hiring less dependent on the availability of unemployed workers, hence more brisk in late stages of expansions. In contrast, when the economy expands and unemployment falls, small firms cannot keep pace because they find it hard to keep hiring. When the economy enters a downturn, large firms have more employment accumulated through poaching that they now want to shed. Small firms were previously more constrained in their growth by search and hiring frictions, and thus now shrink not as quickly.

We contribute to the above literature in several important dimensions. Homebase data provide a rare opportunity to correct for dynamic employee composition effects. The granularity of the data allows us to differentiate between existing workers, new joiners, and quitters. We can thus fix the employee composition at horizon 0 for each period and then trace out the impact for that same group of people over the length of the impulse response function (we also apply the same composition for firms where appropriate). Unlike Moscarini and Postel-Vinay (2012), we are able to analyze the dynamic responses of employment over time. Additionally, we employ a novel instrumental variable approach to ensure that we get true exogenous variation, using detailed USASpending data. These data so far have been used in different contexts, including to analyze local government multipliers (Auerbach et al., 2020a).

A related strand of the literature investigates the importance of employment-to-employment transitions in determining wages. Under certain restrictions, the rate of job-to-job (JJ) transitions is, adjusting for productivity, nearly a sufficient statistic for the average wage (Moscarini and Postel-Vinay, 2016b). Empirical evidence confirms that wages are substantially more sensitive to changes in the arrival rate of job offers to the employed than to changes in the arrival rate of job opportunities to the unemployed, in contrast to predictions of traditional search models (Moscarini and Postel-Vinay, 2016b; Karahan et al., 2017). Moscarini and Postel-Vinay (2016b) explain this with (i) a composition effect: Workers typically quit a job when they receive a better offer, hence the faster these transitions the higher the pace of reallocation toward high wages, and (ii) a strategic effect: the more opportunities workers have to quit, the more aggressive are their employers with their wage responses, to try and retain them. While composition effects benefit only job movers, strategic effects benefit both movers and stayers. Therefore, wage growth is positively related to the pace of JJ reallocations for all workers, but especially for movers. Moscarini and Postel-Vinay (2017) verify this empirically with longitudinal micro data from the Survey of Income and Program Participation. More recently, Moscarini and Postel-Vinay (2023) develop a model which introduces on-the-job search frictions in an otherwise standard monetary DSGE New-Keynesian model. They show that competition for employed, not (only) unemployed, workers transmits aggregate shocks to wages. Since we are able to observe the dynamic wage response within firms over time, differentiating between stayers and new hires our paper contributes to this second strand of literature.

Our paper is also related to a growing literature that exploits payroll data. Previous studies using the Homebase data have focused on analyzing the economic impacts of COVID-19 in real time (Chetty et al., 2023) and studying the pandemic's initial labor market impact including on employment patterns and hours worked (Bartik et al., 2020). The dataset has also been used to investigate the impact of certain pandemic-era policy interventions, including the Paycheck Protection Program (Granja et al., 2022; Kurmann et al., 2023) and the expanded unemployment insurance programs (Altonji et al., 2020). Our paper differs from these earlier studies in important aspects. Most of the earlier analyses use the publicly available Homebase data while ours uses the proprietary data with wages and other job-level information not available in the public data. This allows us to examine wage outcomes in addition to the employment and business closure outcomes. It also allows us to examine a much longer time period including pre- and post-COVID-19, as opposed to the short period immediately after the onset of COVID-19. We thus add to the literature by analyzing the responses to business cycle and labor market fluctuations more generally. This is especially important to understand the evolution of employment and wages across different types of firms in the post-pandemic tight labor market. A contemporaneous paper also uses wage information from the proprietary Homebase data but unlike ours, it focuses on the persistent effect of COVID-19 shock on post-pandemic wage and earning dynamics (Chen and Lee, 2024).

The remainder of this paper is organized as follows: Section 2 introduces our data, Section 3 lays out the methodology and results of OLS regressions and explains how we correct for dynamic composition changes. Section 4 discusses our instrumental variable approach using detailed USASpending data. Section 5 concludes.

2 Data

We are interested in the relationship between local labor market conditions and outcomes for small firms. We therefore construct two sets of variables, one for each side of this relationship. The first is a measure of local labor market tightness, namely the vacancy-to-unemployment ratio. The second consists of measures of employment, wages, hours, and firm survival aggregated from individual firms. Given data availability, our unit of observation for the series we construct is the county-month, and the period of observation is January 2018-December 2022. We also construct a third variable, local government spending, aggregated from data on government procurement contracts, but since we use it only in our instrumental variable analysis we defer discussion to Section 4.

2.1 Local labor market tightness

The vacancy-unemployment ratio as a measure of labor market tightness has its origins in search models of frictional unemployment, initially developed in Blanchard and Diamond (1989). This theory argues that the ratio of firms looking for workers (i.e. vacancies) to workers looking for jobs (unemployment) is a sufficient statistic for the state of the labor market.[5] As a result, this ratio is commonly used in empirical applications as a summary measure of labor market “tightness”.[6]

For our analysis, we compute the ratio of local vacancies to local unemployment. The latter is straightforward to collect, with county-level monthly unemployment published by the Bureau of Labor Statistics as part of the Local Area Unemployment Statistics. To compute a measure of local vacancies, we use proprietary data from Indeed. Indeed is a worldwide search engine for job listings. The Indeed dataset collects job postings anywhere on the internet including job listing sites, employer career sites, and applicant tracking

[5] See Pissarides (2000) for a textbook treatment.

[6] Recent examples include Blanchard et al. (2022), who discuss this measure in the context of the post-COVID U.S. labor market, and Duval et al. (2022), who use the vacancy-unemployment ratio to measure labor market tightness in an international setting.

systems. Duplicated listings are removed so that when the same job is collected from multiple sources it is shown only once. The dataset includes 142 million job postings in our final dataset [7] The dataset covers 421 occupations based on ISCO-08 classifications in 2.9 million companies and 576 counties. Detailed summary statistics are shown in Table 1. To attain a proxy for local vacancies, we then aggregate job postings from Indeed by month at the county level. A priori, there are a number of concerns about the applicability of the Indeed data as a more general proxy for job openings.[8] Nevertheless, as Figure 2a shows, total jobs on Indeed display similar dynamics to total job openings from the official JOLTS statistics at a national level. Figure 2b shows that the correlation between the two statistics at the state-month level is as high as 0.96.

One concern about this measure of labor market tightness is that county-level vacancies may not be terribly informative in an era of remote work – something particularly pertinent in the period we study. Although we cannot rule out this concern entirely, the firms that we study are unlikely to be particularly badly affected by this issue, since they are principally in service industries (especially food services and retail) which are not naturally well-suited to remote work (see Table 3).

year	Job postings	Occupations	Companies	Counties
2019	27,691,685	421	1,508,218	576
2020	25,900,182	421	1,180,153	576
2021	43,214,041	420	1,345,250	576
2022	44,900,736	420	1,163,524	576
Whole sample	141,706,597	421	2,867,705	576

Table 1: Summary statistics: Indeed job postings

Table 1 shows summary statistics for the final Indeed dataset that is used to generate vacancies by county and month, an input to the vacancy-to-unemployment ratio. "Job postings", "Occupations", "Companies" and "Counties" are the total numbers of job postings, occupations (as classified by ISCO08), companies and counties respectively. The number of job postings in the whole sample (last row in the table) is slightly below the sum of the yearly job postings, as the same job posting appears in multiple years if listed in late December and continuing on to January.

[7] The sample period starts in January 2019 and ends in December 2022, same as the sample period for the empirical analysis. The sample period for the outcome variables in Homebase starts a year earlier to allow for the calculation of annual changes.

[8] For example, job postings on Indeed do not reflect a precise number of available jobs, as an opening may remain online for a period of time after being filled, or may not be advertised online at all. And shifts in remote work during the pandemic may further undermine the usefulness of this measure.

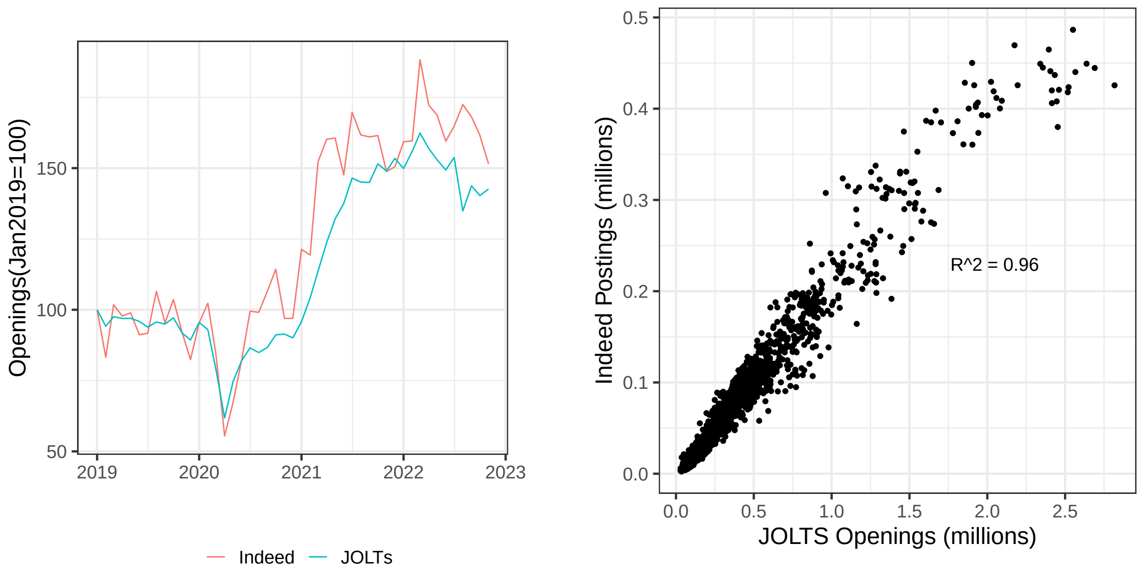

(a) Vacancies: Indeed v.s. JOLTS

(b) Indeed Postings and JOLTS Openings, by state

Figure 2: Job Postings on Indeed vs. Openings on JOLTS

Figure 2a compares the monthly national aggregate job postings on Indeed with the monthly national aggregate job openings on the Job Openings and Turnover Survey (JOLTs) from the BLS. Data for the Figure are scaled such that the level in January 2019 equals to 100. Figure 2b is a scatter plot that compares the monthly number of job postings on Indeed with the number of job openings reported by JOLTS in the same state.

2.2 Outcomes for small firms

Our measures of outcomes for small firms are constructed from raw data provided by Homebase. Homebase is a software company that provides scheduling, payroll reporting, and other services to small businesses, primarily in the retail, hospitality, and other services sectors. This proprietary dataset is based on timecard records provides work hours and wages for over 80,000 businesses and over 1 million employees in the U.S. The data consist of daily records of employee-level information on hours, wages, and job-level information on the duration and the type of jobs (manager or non-manager positions). Employee-level records are linked to the establishment where they work and the firm that owns the establishment. Detailed information on location (5-digit zip code) and industry (6-digit NAICS code) classification are provided at the establishment level. The sample coverage in our study is from January 2018 to December 2022. The number of firms and workers in our sample grows slightly over time (with 2020 an obvious exception, see Table 2).

The dataset's rich coverage of small and private businesses in the retail trade, education and health, leisure and hospitality, and other services and the level of detail on work hours and wages is unique, compared to other micro datasets for the U.S. labor market (see Table 3). Commonly used firm- or establishment-level micro datasets such as Compustat and census lack wage information for small and private firms. This focus on certain industries is not without costs, though; it is not representative of aggregate employment. It also lacks information on tips, benefits and overtime payments. Despite these limitations, employment data from the Homebase sample shows a strong correlation with official statistics. The national employment and earning trends observed in the Homebase sample closely aligns with that calculated from the Current Population Survey (CPS) and the Current Employment Statistics survey (CES). Additionally, changes in employment and earning observed in the Homebase sample show a strong correlation with the CES at the month-state level (Dvorkin and Isaacson, 2022; Chen and Lee, 2024).[9]. Moreover, the dataset is particularly useful for studying the low-wage, low-skill in-person services sector.

Year	Counties	Workers	Firms	Avg. firm size	Workers pc (%)	Firms pc (%)
2019	576	2,267,647	276,534	8.1	0.07	0.009
2020	576	2,099,326	275,913	7.6	0.07	0.009
2021	576	2,897,678	366,362	7.7	0.10	0.012
2022	576	3,295,358	398,834	8.1	0.12	0.015
Whole sample	576	10,560,009	1,317,643	8.0	0.09	0.011

Table 2: Summary statistics: observations, by year

Table 2 shows summary statistics for the number of observations in the final dataset, split by location. “Counties” is the number of counties with sufficient population and a full set of 47 observations. “Workers” and “Firms” are the total number of worker-month and firm-month observations respectively. “Avg. firm size” is the average firm size across counties and periods within states, and so may differ slightly from the ratio of Workers to firms. “Workers pc” and “Firms pc” express respectively the number of workers and firms per capita, and so give a sense of the variation in relative coverage.

[9]The CPS, co-sponsored by the Census Bureau and the BLS, surveys about 60,000 U.S. households and serves as the primary source of official unemployment statistics. The CES, sponsored by the BLS, surveys about 145,000 U.S. businesses and government agencies and serves as the primary source of official employment and wage statistics.

Industry	Workers	Percent	Firms	Percent	Avg. firm size
Full-Service Restaurants	2,804,779	26.6	267,698	20.3	9.5
Unknown	1,470,023	13.9	173,672	13.2	7.4
Other	1,279,922	12.1	209,718	15.9	5.4
Limited-Service Restaurants	1,116,434	10.6	122,187	9.3	8.2
Snack and Nonalcoholic Beverage Bars	769,221	7.3	86,510	6.6	8.1
Drinking Places (Alcoholic Beverages)	453,363	4.3	46,121	3.5	8.6
Women's Clothing Stores	227,271	2.2	42,605	3.2	4.7
All Other General Merchandise Stores	209,500	2.0	29,442	2.2	6.6
Offices of Physicians (except Mental Health Specialists)	198,081	1.9	36,266	2.8	5.0
Museums	177,822	1.7	16,212	1.2	9.6
Convenience Stores	165,781	1.6	25,716	2.0	5.8
Corporate, Subsidiary, and Regional Managing Offices	161,090	1.5	25,720	2.0	5.5
Supermarkets and Other Grocery Retailers (except Convenience Retailers)	149,039	1.4	17,599	1.3	7.7
All Other Amusement and Recreation Industries	145,297	1.4	13,593	1.0	9.4
Elementary and Secondary Schools	140,039	1.3	15,654	1.2	7.5
General Automotive Repair	105,072	1.0	22,162	1.7	4.3
Marketing Consulting Services	87,464	0.8	13,103	1.0	5.4
Other Personal Care Services	81,350	0.8	16,211	1.2	4.4
Furniture Stores	76,211	0.7	16,298	1.2	4.2
Beauty Salons	75,126	0.7	18,643	1.4	3.6
Child Care Services	70,580	0.7	9,577	0.7	6.5
Tobacco Stores	70,409	0.7	16,773	1.3	3.9
Hobby, Toy, and Game Stores	70,173	0.7	15,040	1.1	4.1
Exam Preparation and Tutoring	67,816	0.6	9,348	0.7	6.4
Fitness and Recreational Sports Centers	66,382	0.6	9,958	0.8	5.7
Caterers	59,389	0.6	5,943	0.5	8.4
Services for the Elderly and Persons with Disabilities	55,124	0.5	7,013	0.5	6.9
Pet Care (except Veterinary) Services	54,296	0.5	5,342	0.4	9.0
Florists	52,537	0.5	9,297	0.7	5.0
All Other Miscellaneous Store Retailers	50,382	0.5	8,337	0.6	5.5
Veterinary Services	50,036	0.5	5,885	0.4	7.9
Whole sample	10,560,009	100.0	1,317,643	100.0	7.2

Table 3: Summary statistics: observations, by industry

Table 3 shows summary statistics for the number of observations in the final dataset, split by location. “Workers” and “Firms” are the total number of worker-month and firm-month observations respectively. “Avg. firm size” is the average firm size across counties and periods within industries, and so may differ slightly from the ratio of workers to firms. The two columns labelled “Percent” express the number of workers and firms by industry as a fraction of the total .

We use the Homebase data to build county-level average outcomes for employees per firm, hours per worker, wages per hour, number of firms, and the total number of jobs (see summary statistics in Table 5). The firms in our dataset are overwhelmingly very small; the median firm in our data has fewer than 7 employees, paying an average hourly wage of around \$12. Workers are almost exclusively casual or part-time employees. Some work regularly but other are much more intermittent – with hours on a handful of days one week and none the next. This suggests that the firms we study have highly volatile businesses and offer jobs that are low paid and unreliable. A key challenge in constructing county-level is dynamic sample composition bias. This occurs when there is correlation between the shock of interest and the subsequent composition of the sample. We expect this to be a non-trivial problem in our investigation. The jobs ladder framework of Moscarini and Postel-Vinay (2016a) provides some motivation for this concern. In their model, better-paid, higher-skilled workers have better outside options, and so are first to move to better-paying jobs at larger firms when labor markets are tight. The resulting composition effect can result in lower average wages at small firms even if wages increase for remaining employees. As such, the response of the simple average wage to changes in labor market tightness will confound the change in like-for-like wages with this composition effect.[10]

To address this problem, we develop two corrections which we apply to the dynamic aggregation of the Homebase data. These corrections are better understood in the context of our preferred econometric framework and are discussed in detail in Section 3.1.

3 Estimation via Ordinary Least Squares

3.1 Methodology

Our first analysis of the data is via ordinary least squares estimation of local projections, following Jordà (2005). Specifically, for horizons $h = 1, \ldots, H$ we estimate:

$$y^h_{c,t} - y^0_{c,t} = \alpha^h_t + \eta^h_c + \beta^h \theta_{c,t} + \sum_{k=1}^{K} \gamma^h_k \theta_{c,t-k} + \sum_{m=0}^{M} \delta^h_m y_{c,t-m} + u^h_{c,t} \tag{1}$$

where observations are indexed by county c and month t, $y^h_{c,t} - y^0_{c,t}$ is the difference in the dependent variable of interest over a horizon of length h starting at time t, α^h_t is a time fixed effect, η^h_c a county fixed effect, and $\theta_{c,t}$ is our independent variable of interest, local labor market tightness. We also allow for the inclusion of lags of both the dependent and independent variables.

[10]The problem is not limited to wages. This is a general issue stemming from the dynamic change in the composition of the sample at different horizons.

State	Counties	Workers	Firms	Avg. firm size	Workers pc (%)	Firms pc (%)
Alabama	12	116,899	16,331	7.2	0.09	0.012
Alaska	1	6,001	911	6.5	0.12	0.018
Arizona	8	257,529	31,110	8.3	0.08	0.010
Arkansas	7	62,076	7,648	8.1	0.10	0.012
California	34	2,136,421	232,065	9.1	0.12	0.013
Colorado	11	241,134	32,013	7.5	0.11	0.014
Connecticut	7	81,171	11,045	7.4	0.05	0.007
Delaware	3	31,507	3,999	7.9	0.07	0.009
District of Columbia	1	40,625	4,586	8.8	0.12	0.014
Florida	36	925,784	128,423	7.2	0.10	0.014
Georgia	25	380,340	52,376	7.2	0.11	0.016
Hawaii	2	50,397	5,499	9.0	0.09	0.010
Idaho	4	72,378	8,125	8.8	0.16	0.018
Illinois	19	301,088	38,577	7.8	0.06	0.008
Indiana	17	126,400	16,425	7.7	0.06	0.008
Iowa	6	47,780	6,015	7.8	0.08	0.010
Kansas	5	80,291	9,533	8.4	0.11	0.013
Kentucky	7	60,674	7,162	8.4	0.07	0.009
Louisiana	14	161,330	20,403	8.0	0.11	0.014
Maine	4	20,748	3,099	6.7	0.06	0.009
Maryland	14	147,201	21,536	6.8	0.06	0.008
Massachusetts	11	158,511	19,499	8.1	0.05	0.006
Michigan	19	230,497	29,412	7.8	0.06	0.008
Minnesota	11	135,670	15,967	8.4	0.08	0.009
Mississippi	6	46,715	5,819	8.1	0.10	0.012
Missouri	12	169,923	21,274	7.9	0.10	0.012
Montana	2	13,584	2,366	5.8	0.10	0.018
Nebraska	3	76,922	8,743	8.7	0.15	0.017
Nevada	2	142,719	15,943	8.8	0.11	0.013
New Hampshire	3	38,645	3,913	10.0	0.10	0.010
New Jersey	19	228,303	27,666	8.2	0.06	0.007
New Mexico	5	57,042	7,291	7.7	0.09	0.012
New York	26	436,510	50,013	8.7	0.05	0.006
North Carolina	25	375,100	45,755	8.2	0.11	0.014
North Dakota	1	6,710	915	7.3	0.08	0.011
Ohio	24	292,624	38,435	7.6	0.07	0.010
Oklahoma	5	95,537	12,916	7.3	0.10	0.014
Oregon	9	211,918	26,448	8.0	0.14	0.018
Pennsylvania	30	308,703	39,606	7.8	0.06	0.008
Rhode Island	3	18,209	2,499	7.3	0.04	0.006
South Carolina	15	175,664	23,985	7.3	0.10	0.013
South Dakota	2	11,984	1,768	6.7	0.08	0.012
Tennessee	13	194,567	24,717	7.8	0.10	0.013
Texas	40	990,211	136,073	7.2	0.09	0.012
Utah	6	179,896	20,663	8.7	0.14	0.016
Vermont	1	3,978	606	7.1	0.05	0.008
Virginia	16	168,476	22,994	7.3	0.07	0.010
Washington	12	287,243	38,137	7.5	0.10	0.013
West Virginia	3	16,990	2,002	8.5	0.09	0.011
Wisconsin	15	139,384	15,337	9.1	0.08	0.009
Whole sample	576	10,560,009	1,317,643	8.0	0.09	0.011

Table 4: Summary statistics: observations, by location

Table 4 shows summary statistics for the number of observations in the final dataset, split by location. "Counties" is the number of counties with sufficient population and a full set of 47 observations. "Workers" and "Firms" are the total number of worker-month and firm-month observations respectively. "Avg. firm size" is the average firm size across counties and periods within states, and so may differ slightly from the ratio of Workers to firms. "Workers pc" and "Firms pc" express respectively the number of workers and firms per capita, and so give a sense of the variation in relative coverage. NB: These counties have a combined population of around 256 million.

Variable	Year	5%	10%	25%	50%	75%	90%	95%
Avg. firm size	2019	3.9	4.4	5.5	6.6	8.0	9.5	10.9
	2020	3.6	4.2	5.0	6.0	7.2	8.6	9.8
	2021	4.1	4.7	5.5	6.5	7.6	9.1	10.0
	2022	4.4	4.9	5.9	6.9	8.1	9.4	10.4
Avg. hourly wage (USD)	2019	8.24	8.68	9.49	10.47	11.73	13.02	13.66
	2020	8.83	9.35	10.25	11.28	12.77	14.23	15.06
	2021	9.28	9.86	10.90	12.00	13.58	15.04	15.67
	2022	10.25	10.80	11.83	12.98	14.55	15.93	16.58
Avg. weekly hours	2019	18.6	19.9	22.0	24.0	26.0	28.1	29.6
	2020	19.2	20.5	22.5	24.5	26.7	28.9	30.5
	2021	20.0	21.2	23.1	25.0	27.0	28.8	30.0
	2022	20.8	22.0	24.0	26.2	28.6	31.0	32.6
Proxy v/u ratio	2019	0.17	0.23	0.32	0.44	0.61	0.83	0.99
	2020	0.06	0.08	0.14	0.24	0.38	0.55	0.67
	2021	0.21	0.27	0.39	0.55	0.78	1.08	1.30
	2022	0.35	0.43	0.59	0.80	1.12	1.49	1.74
Unemp. rate (percent)	2019	2.3	2.5	2.9	3.4	4.2	5.1	5.7
	2020	3.0	3.5	4.6	6.6	9.9	13.3	15.5
	2021	2.5	2.9	3.7	4.7	6.0	7.4	8.4
	2022	2.1	2.3	2.8	3.4	4.1	5.1	5.8

Table 5: Summary statistics: key series, by year

Table 5 shows summary statistics for the distribution of key variables in the data across counties split by year. Average firm size is workers per firm, Average weekly hours are per worker, and “Proxy v/u ratio” is given by indeed vacancies divided by BLS unemployment rate.

From a sequence of these projections, we are able to trace out the impulse response of the dependent variable to a change in local labor market tightness. We focus on three challenges in interpreting these responses: (i) dynamic composition bias, where the composition of the sample changes endogenously over the duration of the impulse, h; (ii) static composition bias,[11] the risk that the sample at $h = 0$ is different across time periods t in ways correlated with the independent variable, $\theta_{c,t}$; and (iii) questions about the source of the shock.

Dynamic composition bias. A key challenge to interpreting the impulse responses is the possibility that the sample composition changes endogenously over time. One way to correct for this is to use extensive controls in the regression. For example, if more educated workers were more likely to leave firms after an increase in labor market tightness, then including the average level of education of observed workers in period $t+h$ in the horizon-h regression would correct for this. However, even with the most extensive set of controls, it is impossible to be certain that they are sufficient to address the problem.

Because we have individual-level data on worker-firm matches, we can correct the data directly at source to avoid dynamic composition bias without controls. We simply construct aggregate county-level measures for wages, hours, employment, firm survival, and worker-firm match survival by compare *exactly the same people* (or firms) at each horizon h following any given period t.

In defining the composition correction calculations, we focus on average wages. For the other four variables, the details are slightly different but the principles are exactly the same. We discuss some of the differences in Section 3.2. Let $E_{c,t}$ be the sample of workers in county c in period t. Worker j works l_t^j hours in period t at wage w_t^j. One can then construct three measures of the firms' average labor cost per hour h periods after period t. These are:

$$\text{Sample average:} \qquad \bar{w}^h_{c,t} = \frac{\sum_{j \in E_{c,t+h}} l_t^j w_t^j}{\sum_{j \in E_{c,t+h}} l_t^j}$$

$$\text{Entry-corrected:} \qquad \tilde{w}^h_{c,t} = \frac{\sum_{j \in E_{c,t}} l_t^j w_t^j}{\sum_{j \in E_{c,t}} l_t^j}$$

$$\text{Double-corrected:} \qquad \hat{w}^h_{c,t} = \frac{\sum_{j \in \bigcap_{k=1}^{H} E_{c,t+k}} l_t^j w_t^j}{\sum_{j \in \bigcap_{k=1}^{H} E_{c,t+k}} l_t^j}$$

All three versions calculate some measure of total wages and divide by total hours. The difference between them is which workers are included. The first is the simple sample

[11] The static and dynamic composition problems are both forms of selection bias. Each stems from certain workers or firms selecting into our sample in ways correlated with the independent variable. However, the static composition bias seems a more standard form, with more standard remedies.

average, which just aggregates over all workers in each period. This is a benchmark and does not adjust for sample composition change. The second, the entry-corrected measure, fixes the sample at only the workers employed in period t. This means that any workers joining the firm in periods $t+1\dots,t+H$ are not included in this measure. However, workers who leave will drop out of the sample. The third measure, the double-corrected one, restricts the sample only to those that are observed over the whole horizon. This is a fully consistent sample, comparing like with like over the horizon of the impulse response and uncontaminated by entry or exit of workers from our pool of observations.

Table 6 provides a simple example. Here, we consider a hypothetical county observed for six periods, where four workers–A, B, C, and D–work intermittently, with their appearance in the full sample described in the top two rows of Table 6.

Period, t	1	2	3	4	5	6
Workers, $E_{c,t}$	A	A,B	A,B,C	A,B,C,D	B,C,D	C,D
<u>$t = 1$</u>						
h	0	1	2	3		
Uncorrected	A	A,B	A,B,C	A,B,C,D		
Entry-corrected	A	A	A	A		
Double-corrected	A	A	A	A		
<u>$t = 2$</u>						
h		0	1	2	3	
Uncorrected		A,B	A,B,C	A,B,C,D	B,C,D	
Entry-corrected		A,B	A,B	A,B	B	
Double-corrected		B	B	B	B	
<u>$t = 3$</u>						
h			0	1	2	3
Uncorrected			A,B,C	A,B,C,D	B,C,D	C,D
Entry-corrected			A,B,C	A,B,C	B,C	C
Double-corrected			C	C	C	C

Table 6: Sample correction example

Table 6 shows the the three sample correction methods employed for a hypothetical example where workers A, B, C, and D are observed in periods $t = 1, \dots, 6$ with a maximum horizon length of $H = 3$. The top two rows show the raw samples. The bottom half shows the samples used for the three sample-correction methods we use in constructing the observations for each of $t = 1, 2, 3$

Table 6 illustrates the dynamic composition problem. Using uncorrected data in the $h = 1$ regression, the $t = 1$ observation will be the average of wages of workers A and B. But in the $h = 2$ regression, the $t = 1$ observation will be the average over workers A, B, and C. And at horizon $h = 3$ the sample composition is different yet again. A similar problem applies to the observations for $t = 2$ and $t = 3$. As such, the regression coefficients β^h will trace out an impulse response which does not capture the response of wages for a given type of labor, but instead the change in the average wage of the employed. This combines something economically meaningful–the price of a unit of labor–with something less so–the composition effect. In large samples, this is not necessarily a problem; if there is no correlation between the independent variable and the sample composition, then with enough data the composition effect will wash out and one will be left with an estimate of the change in the wage. However, in our case, the job-ladder theory predicts that there should be exactly this correlation. If higher paid workers leave to better jobs first in a tight labor market, then the composition effect will bias downwards the estimated response of wages.

To address this concern, the two corrected versions fix the sample, either to include only those observations available at time t (the entry-corrected version) or only those which are available throughout the impulse (the double-corrected) version. At the example in Table 6 should make clear, only the latter version fully addresses the dynamic composition issue, comparing like with like across the different horizons (the sample still varies *across* periods–this is static composition bias, which we discuss below). The entry-corrected version has some value for computing statistics based on counts of observations, such as the fraction of firms or worker-firm matches which persist h periods after a shock. There, the double-corrected measures will exhibit zero variation by construction – only the entry-corrected statistics (which fix the sample at the time of the shock) are meaningful.

The same correction techniques can be applied when the unit of observation is the firm. For example, when computing the number of employees per firm, one may want to distinguish between variation due to the number of workers per firm for a constant set of firms versus that due to changes in the composition of the sample of firms. To calculate the former, one would apply composition correction to firms. Finally, it is important to note that none of the sample correction techniques is inherently better or worse than the others, only different. Which one is appropriate depends on the interpretation one wishes to give the estimates.

Static composition bias. The sample composition may also change across periods in ways correlated with the independent variable ("static composition bias"). This shows up in the example in the difference in samples across $t = 1, 2$, and 3. Even after correcting for dynamic composition effects, the double-corrected sample in Table 6 changes from A to B to C over time. Note that this is not a problem arising due to the composition correction; it is a feature of the uncorrected sample as well. The only difference is that the double-corrected one provides a consistent sample over the duration of the impulse.

We take two steps to address static composition bias. First, we include time fixed effects as well as lagged dependent and independent variables. To the extent that changes in the set of workers or firms is correlated with these factors, then this should produce estimates of the impulse response which adjust for sample composition changes across the business cycle. Second, we take advantage of the ability to condition local projections on other covariates to compute separate marginal impulse responses at different levels of labor market tightness.

The source of the shock. One final issue with interpreting the impulse responses is in understanding what drives the impulse response. Typically, empirical studies try to isolate exogenous variation in the independent variable. In our case, however, it is not obvious what exogenous variation labor market tightness might even mean. The vacancy-to-unemployment rate is an endogenous outcome, driven by macro shocks. One interpretation is that this is a demand shock, on the grounds that only local demand shocks are likely to cause fluctuations in residual variation in local labor market tightness at monthly frequency after accounting for common time effects and lagged variables. In Section 4 we offer a more robust defense of this interpretation when we present estimates calculated by instrumental variables, where we can more plausibly claim that the variation is due to locally exogenous demand shocks.

3.2 Headline results

We estimate equation (1) using a county-month dataset constructed from our underlying data sources. In our baseline we use three lags of each of the dependent and independent variables (so $K = M = 3$) and in Section 3.3 we verify that our results hold for alternate lag structures. One risk with is that because less-populated counties will typically have fewer observations in the micro-data, the constructed measures from them will be more noisy. In larger counties the law of large numbers is more likely to apply. To reduce this possibility, we weigh observations by county population and omit counties with populations smaller than 100,000. Given that our data may underrepresent primary industries, we also omit the most rural counties, including only counties in categories 1-4 of the CDC's 6-part urban-rural classification ("small metro" and larger). This produces a balanced panel of 576 counties with observations in every month from January 2018 to December 2022. We relax these assumptions in the robustness exercises presented in Section 3.3.

Figure 3 presents our headline results using ordinary least squares. Dotted lines show a 95 percent confidence interval using standard errors clustered by month and county. The first panel shows the dynamic auto-response of local labor market tightness. On impact, labor market tightness increases one percentage point, and decays relatively quickly; it is statistically indistinguishable from zero after 6 months.

Following a one percentage point increase in labor market tightness, labor inputs *contract* on both the extensive and intensive margins. After 6 months, the number of employees per firm has fallen by about 0.1 percent and the average hours per employee by around 0.04 percent (panels 2 and 3 of Figure 3 respectively). Importantly, dependent variables in panels 2 and 3 are double corrected by firm and worker respectively. It thus rules out a simple explanation for our results, namely that they are purely a result of the changing composition, for instance larger firms or workers with more hours dropping out of our sample. Instead, our results suggest that when labor markets are tight, a given pool of small firms sheds workers and gives a fixed set of workers fewer hours.

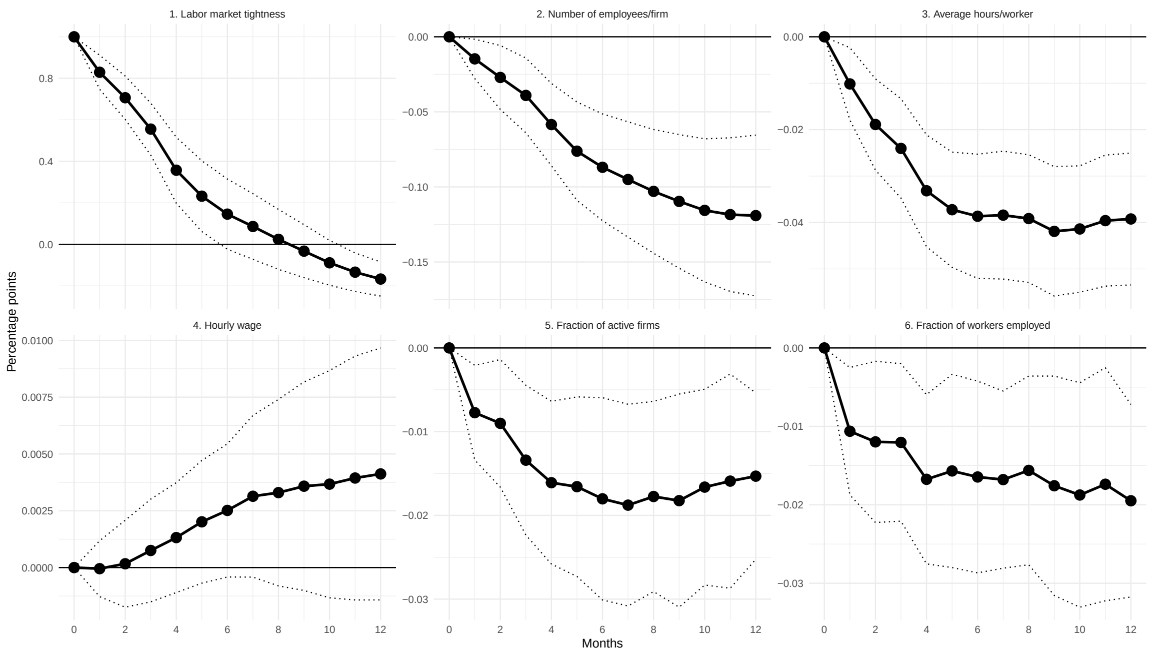

Figure 3: Headline results: Ordinary Least Squares

Panel shows responses of outcome variables to a one percent increase in labor market tightness. Dotted lines show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations.

Wages also appear to increase following the tightening shock, although the quantitative and statistical significance is less clear (see Panel 3 of figure 3). The six-month elasticity to a one percent tightening in local labor markets is only around one quarter of one percent. Because we include time fixed effects, we can interpret these as real wage increases, at least when discounted by national price indices. This is consistent with the notion that small firms have to compete for workers in an environment where local demand is strong, pushing up wages. Firm failure and worker separation rates also go up. This can be seen in panels 5 and 6 or Figure 3. These are entry-corrected by firm and worker respectively, and so show the fraction of firms in the data (panel 5) and workers in the sample (panel 6) at period $t+h$ which were observed at period t.[12] Panels 5 and 6 suggest that on average, 0.02 percent more firms (and workers) exit our sample following a one percent increase in labor market tightness. Compared to hours and wages, the interpretation of these last two panels is more ambiguous. It could be that firms are failing and workers are becoming unemployed. Or it could be that firms in distress cut costs, including on payroll software. However, Homebase is not expensive. The cost ranges from free, for a basic version for firms with fewer than 20 employees, up to $80 per month. Thus, the former (firms failing and workers becoming unemployed) is a more likely explanation for our results.

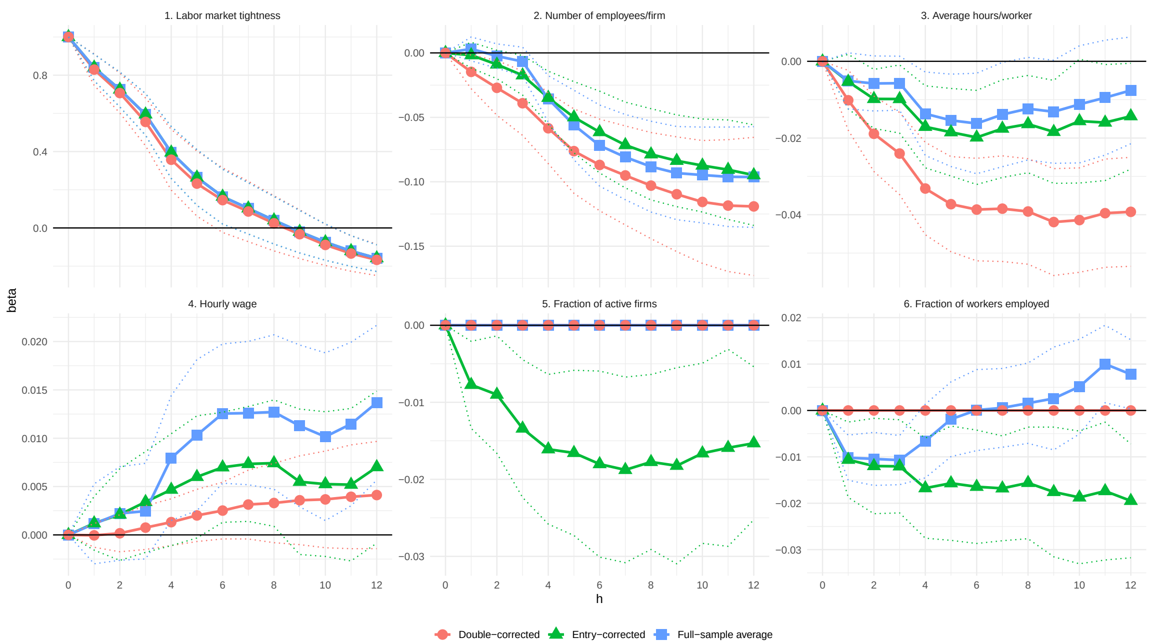

Figure 4: Impact of sample correction: Ordinary Least Squares

Panel shows responses of outcome variables to a one percent increase in labor market tightness. Dotted lines show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Different lines vary by the dynamics sample correction method applied.

In Figure 4 we further examine the role of sample correction. First, in panel 1 we check that the dynamic response of labor market tightness is the same and confirm that the shock is the same independent of how the other dependent variables are constructed. Panel 2 shows

[12] Double-correction does not make sense for these measures. Here, we want to measure whether firm (or worker) attrition is larger or smaller following a period of tight labor markets. Correcting for firm (or worker) exit would remove that variation entirely.

little variation in the total number of employees per firm across measures. However, panel 3 shows that the reduction in hours per worker would be understated without correction: the uncorrected and entry-corrected lines show a much smaller response of hours. This is consistent with departing workers working fewer hours at the time of the shock t. Perhaps less attached to employment, those workers leave first when small firms downsize in tight labor markets. For wages, the uncorrected estimates imply a rather different story. There, the uncorrected sample shows strong growth, driven by entrants. This would be the case if new hires have to be competitive whereas matching frictions may allow firms to keep wages for ongoing employees at below-market rates, at least to some extent.

These findings suggest that controlling for dynamic composition bias is key also in the context of thinking about the slope of the wage Phillips curve. Recent studies have suggested a steepening for low wage workers following the COVID-pandemic (Autor et al., 2023), spurred by transitions to better and higher paying jobs. Autor et al. (2023) find that job-switching is a key predictor of wage growth with job-to-job separation rates exhibiting independent power for predicting cross-state wage growth, reflecting movements along state-level wage-Phillips curves. Our findings shed further light on this work. We confirm that the impact of labor market tightness on wage growth does not seem to be associated with a better bargaining power of workers within a given job (as suggested by traditional models, e.g., Blanchard and Galí (2010)), but rather that wage growth is driven by new entrants, consistent with a job-ladder model.

The differences between the various measures in Figure 4 also have a value beyond the purely statistical. They also tell us something about the employment conditions of those joining and leaving small firms. In panel 4, the gap between the double- and entry-corrected responses implies that those leaving observed firms after labor markets tighten had fewer hours on average than those remaining.[13] A similar gap in panel 4 says that leavers have on average lower wages than remaining workers. The yet higher full-sample line in this panel implies further that incoming workers have higher wages than even those workers who remain throughout. Between them, these responses suggest that there may be some variation in bargaining power amongst workers, which interacts with their decision to leave or to join a firm. Workers with lower bargaining power – for example, if they do not share a family connection with the business owner – may how lower wages and hours and so be more likely to leave when the labor market picks up.

The results in panels 5 and 6 provide perhaps less economic insight but a useful further

[13]This may seem counterintuitive, but because leavers drop out of the sample, they have the opposite effect to what one might expect. In this case, if leavers have fewer hours when they are observed (when h is small), then when they drop out of the sample, the average hours for those remaining would go up. This counters the reduction in hours from those who remain throughout the impulse response (i.e. the double-corrected sample) and leads to an entry-corrected impulse response above the double-corrected one.

cross-checks on our analysis. For panel 5, the impulse responses for both the full-sample and double-corrected measures are zero. This is as it should be – the uncorrected sample includes all firms observed at each point in time which are, by definition, active. Instead, the entry-corrected line here is the most useful – it measures the probability that a firm which was active at time t is still active at time $t+h$. This is why we use single-correction in our headline estimates for this panel. Likewise, the double-corrected sample includes only those firms which do not fail by time $t+h$ and so never varies. For workers, the mechanics are a little different. As with firms, the double-corrected probability of remaining in the sample never changes, because this sample selects only those workers which remain employed throughout. However, the uncorrected sample shows variation because workers often reduce hours to zero for a month or two and then resume work.[14] And so there can be workers who are in the sample but inactive in a given month. The increase in the uncorrected fraction of active workers is thus hard to interpret – it is both a function of the activity of pre-existing workers and new workers. In contrast, the entry-corrected number of workers employed has a more useful interpretation – it is the fraction of workers who were employed at a given firm in time t who are still active at time $t+h$.

In general, the double-corrected measures also have tighter standard errors, especially for wages. This could result if employees who stay at firms over a long period of time have less idiosyncratic variation in their skills or effort.

Together, these results suggest that tight labor markets—which one would usually associate with strong aggregate demand—are tough for small firms. They shed workers and reduce the hours of those remaining, even though costs (probably) go up. More firms fail and more workers leave their jobs. The next section investigates the robustness of these findings.

3.3 Robustness

We conduct four robustness exercises, recomputing our estimates under alternative econometric specifications, for alternate sample restrictions, for different firms sizes, and in different subsamples. Throughout, we present the $h = 6$ horizon results as a summary of our findings, with $h = 1$ and $h = 12$ equivalents in Appendix A.

Figure 5 presents results from five alternate forms of equation (1). The first three consider alternate lag structures, either adding lags of the independent variable, $\theta_{c,t}$, or adding extra lags of both the dependent and independent variables, or removing lagged variables. These give generally very similar results to the baseline specification, with employment and hours falling, wages increasing but there is with less obvious statistical significance on firm

[14] This is in contrast to firms, who usually find that reducing hours to zero is an absorbing state.

and employee attrition. We also consider a specification without time fixed effects. This allows for more variation in the data, but at the potential cost of allowing spurious correlation between the dependent and independent variables. Results under this specification are stronger for wage, but weaker for firm and employee attrition. As a broader check on our work, we also include a placebo experiment here, which replaces the dependent variable with uncorrelated white noise. As expected, the effects throughout are indistinguishable from zero.

Figure 6 shows the impact of some of our sample restrictions, allowing for changes in weights, and adding or subtracting smaller and more rural counties. The only variable for which there is a consistent and noticeable pattern is wages. The increase in wages we estimate is larger for the most urban counties and lowest when we relax all restrictions (i.e. with neither size nor urbanity requirements). This is consistent with the idea that labor markets are deepest in larger, more densely populated areas. So firms there may enjoy a lesser degree of monopsony power, requiring them to raise wages in tight labor markets to remain competitive.

Figure 7 splits the sample by firm size at time t.[15] The results are consistent across firms of different sizes, although uncertainty is larger for firms with the most employees, simply because there are relatively few of them in our sample. Figure 8 also cuts our sample by firm characteristic, in this case for the six most prominent industries in our dataset. Once more, the results hold in general, although with some variation. In general, the departures from the baseline, and the standard errors, tend to be larger for the less-representative industries. Given that the smallest two industry groups only cover around 4 percent of workers between them, it might instead be surprising that there is not more variation in the estimates.

Finally, Figure 9 repeats our estimation on three sub-periods: before, during, and after the COVID-19 pandemic. The results are similar to before, with the qualitative picture remaining similar although uncertainty is larger.

4 Instrumental variables

4.1 Data

There could be several factors that explain changes in labor market tightness, which complicates the interpretation of the results in 3.2. Canonical models of job search show that a number of economic forces, including aggregate demand shocks, revaluation of households'

[15] Another benefit of building the county level outcomes from microdata is that we can account for dynamic composition bias in firm size. That is, we can fix firm size at time t and track the evolution of outcomes for that firm or its employees over subsequent periods without the sample being contaminated by firms shifting size category.

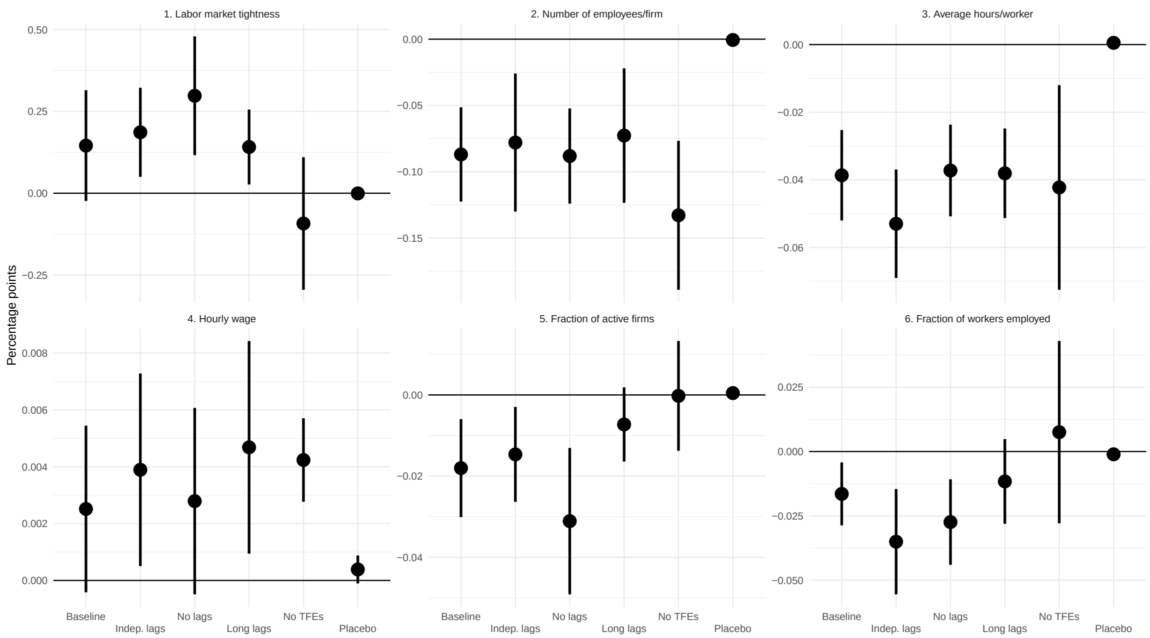

Figure 5: Robustness of OLS estimates: alternate specifications, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Indep. lags* adds 3 lags of the independent variable, *Long lags* includes 12 lags of both the independent and dependent variables, *No lags* omits lags of both dependent and independent variables, *No TFEs* omits time fixed effects, and *Placebo* replaces the independent variable with a white noise process with the same standard deviation as the unemployment-vacancy ratio in the data.

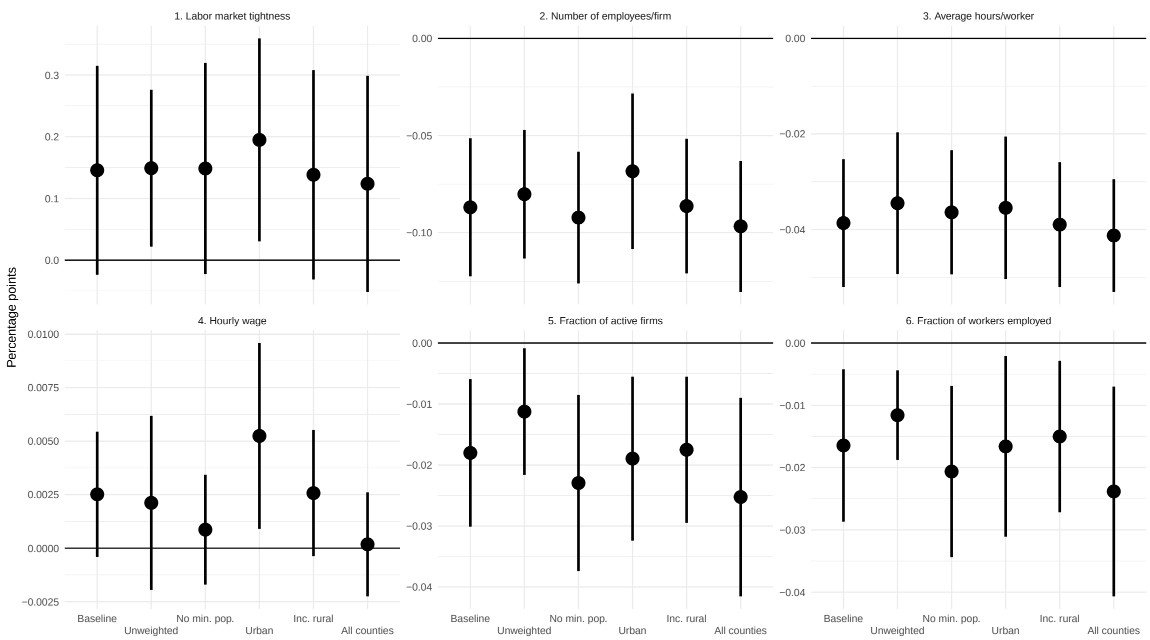

Figure 6: Robustness of OLS estimates: alternate sample restrictions, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Unweighted* removes population weighting, *No min. pop.* removes the minimum county population requirement, *Urban* uses only counties with a CDC Urban score of 1 or 2, *Inc. rural* includes counties with CDC urban score up to 6 (i.e. no restriction), and *All counties* includes all counties without restriction on population or urban/rural character.

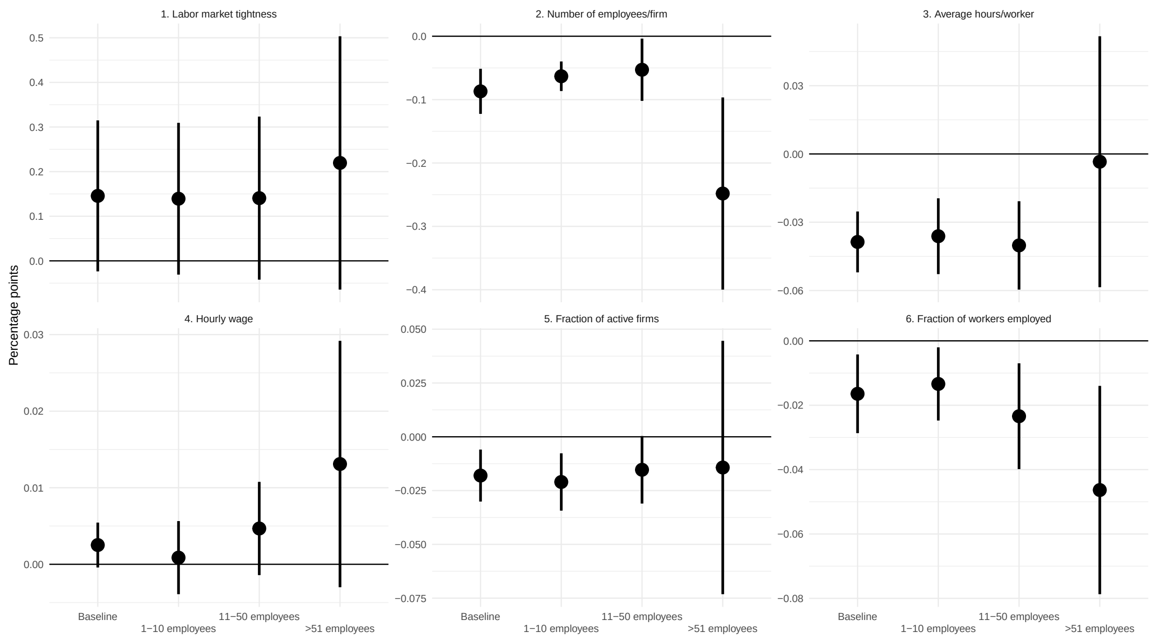

Figure 7: Robustness of OLS estimates: firms size, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. Estimates vary by firm size categories, as indicated in the chart.

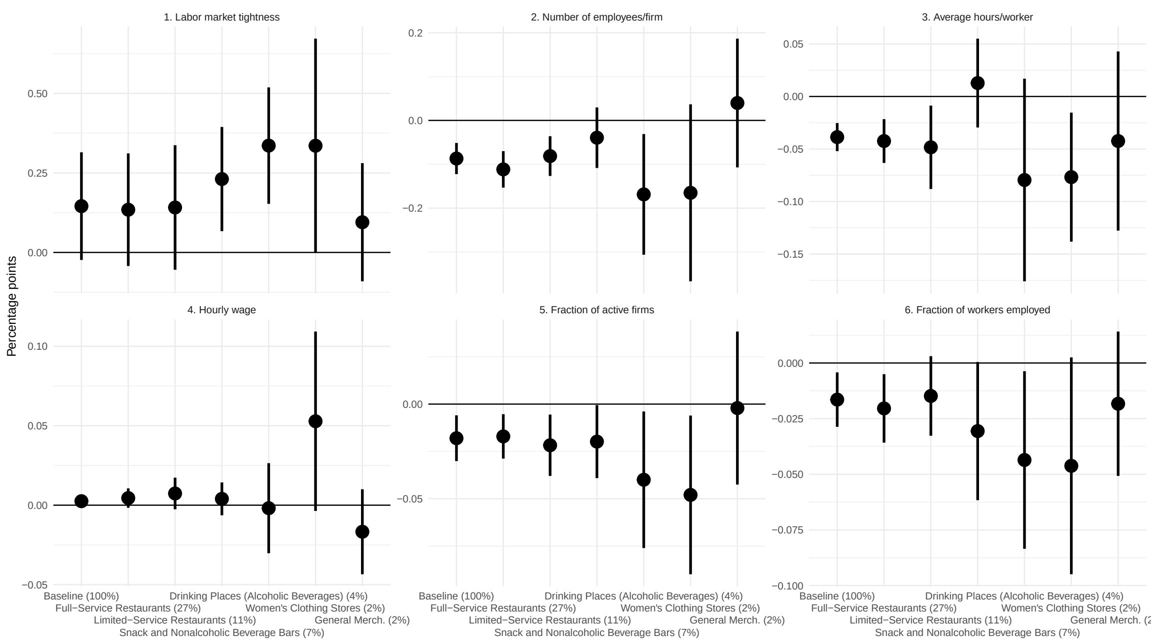

Figure 8: Robustness of OLS estimates: industry-specific subsamples, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Figure shows the 6 most common industries, covering some 53 percent of workers in the sample. Percentages in parentheses are the share of sample workers in that industry.

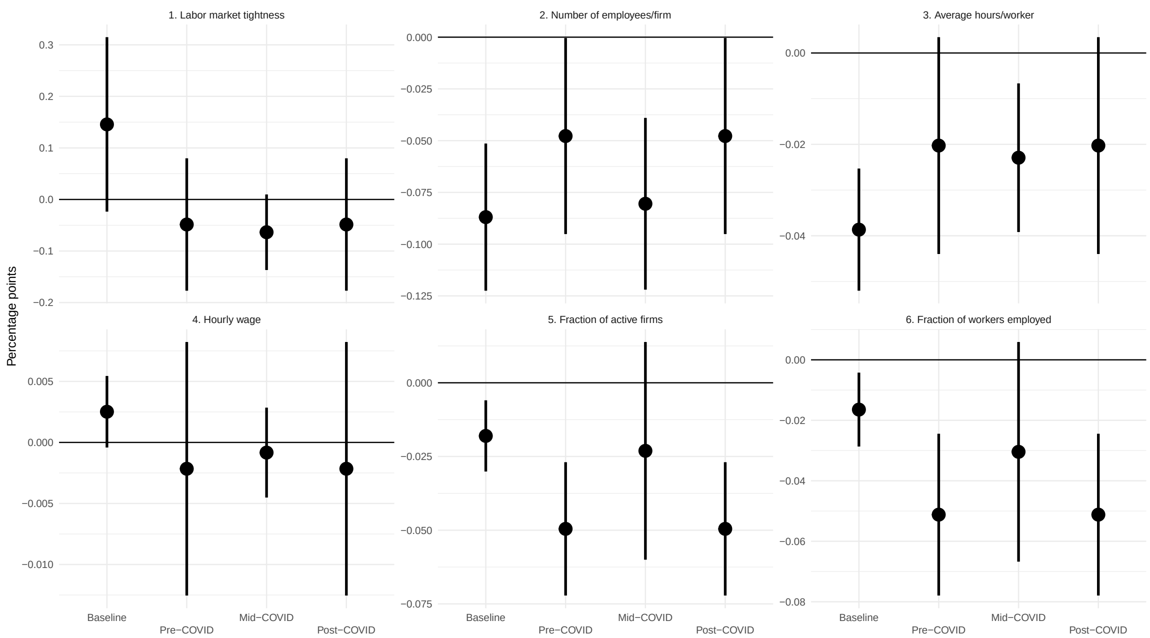

Figure 9: Robustness of OLS estimates: alternate sample periods, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Estimates vary by period, *Pre-COVID* is Jan 2019-Feb 2020, *Mid-COVID* is Mar 2020-Jun 2021, and *Post-COVID* is Jul 2021-Nov 2022.

outside options, as well as changes in matching and production technologies, could drive changes in the vacancy-to-unemployment ratio.[16]

This is problematic since the impact of different shocks on the dependent variables will vary in our regressions, complicating the interpretation of the impulse responses. For instance, an increase in productivity increases both the surplus from existing matches and the value of new vacancies. In equilibrium, both wages and labor market tightness increase. In contrast, an increase in workers' bargaining power lowers the value of new vacancies, causing wages to rise but labor market tightness to fall. Simple regression risks recovering an average impact of all the forces which might induce co-movements in labor market tightness and wages or other variables. Of course, controls for time and county fixed effect and lagged dependent and independent variables can mitigate this concern, but cannot remove it entirely.

To try to address this problem, we construct an instrument for local aggregate demand, aiming to isolate the part of the variation in local labor market tightness stemming from demand shocks. We focus on these shocks given their relevance for policy analysis. Monetary policymakers, in particular, often focus on demand shocks in their analysis of the labor market.[17]

Our instrument is local procurement spending by the United States Department of Defense (DoD). We use transaction-level data from USASpending.gov to create DoD spending by month and county throughout our sample. In this, we follow Auerbach et al. (2020b) in arguing that fluctuations in such spending constitute a shock to local aggregate demand unaffected by local economic conditions. While government spending may, in general, be endogenous to prevailing economic conditions, military spending has long been seen as unrelated, especially at the local level.[18] Moreover, procurement covers purchases of intermediate goods and services from businesses, and omits spending that might be particularly sensitive to local economic conditions.[19] This is in contrast to other forms of spending, such as unemployment benefits, which have a natural and automatic relationship to the economic cycle.

The raw data on USASpending.gov are for individual transactions. After the first transaction of a contract is made, subsequent transactions of the same contract may be made to

[16]See, for example, Pissarides (2000).

[17]To be clear, we do not claim that we identify "exogenous" variation in labor market tightness; θ_t is the ratio of two endogenous variables, so even the notion of exogenous variation is misleading. Instead, we are isolating the variation from a type of shock.

[18]This literature is surveyed extensively in Ramey (2011).

[19]Hooker and Knetter (1997) is an early example of using local military procurement data as an exogenous shock, albeit with state-year resolution. A more recent paper using a similar approach to construct a government spending shock using the USASpending data in Cox et al. (2020). As in that paper, we have true county-month variation which allows us to go beyond Bartik-style approaches which allow only for proportionate variation in the instrument across units.

adjust the amount of funding (either upward or downard). Each observation in the data includes a wide range of information about the transaction, including the start and end date of performance, the total obligation, the awarding agency, a NAICS industry classification for the purchased item, and the primary place of performance of the contract. The granular information make it possible to construct government procurement spending over time by specific government agency in specific industry or location.

One complication is that an obligation is the contracted amount—a promise made by the government to spend funds—instead of the actual spending or outlay. Outlays occur when money is actually paid out and are the relevant fiscal variable for our purpose to measure demand from government purchases. To construct a proxy for outlays, Auerbach et al. (2020b) assume that spending happens evenly over the lifetime of the contract. They therefore divide the total amount of obligation associated with a contract over its duration to proxy for monthly/daily spending. Similarly, we assume that spending happens evenly over the duration of a transaction.[20] The aggregated monthly defense spending constructed from the USASpending.gov dataset broadly matches the national defense spending released in the National Income and Product Accounts (Figure 10).

In total, close to 15 million DoD procurement contracts are open during 2019-2022, the overwhelming majority of which are small—over three quarters of transactions feature a total expenditure of less than $2,500 (Table 7).[21] However, the bulk of the actual spending is comprised of very large contracts (Figure 10), usually either subcontracting for health insurance or for manufacture of advanced military hardware, such as ships or planes.[22] Most contracts have short duration—over three quarters of contracts are completed within two months. While less than five percent of contracts have a duration of more than 1 year, some extend over more than five years (Table 7). Relatedly, between 2019 and 2022, out of the close to $3\frac{3}{4}$ million outstanding defense spending contracts each year, the majority are new contracts that started in the same year while legacy contracts represent less than 10 percent, Table 8). Nevertheless, longer-duration contracts are relatively larger contracts. Of the around $350 billion spending each year, only about one quarter are new contracts. The geographical and industry coverage is broad (including most states, around 490 counties,

[20]Utilizing granular information at the transaction level instead of at the contract level helps capture the timing of the spending more precisely. For instance, if a transaction allows for US$1 million of spending within one year while the transaction is part of a five-year contract, then allocating the US$1 million over the specific year will capture the timing of spending better than allocating the US$1 million over a five-year horizon.

[21]These include, amongst many other: an order for office supplies in Wisconsin for $28 (link), $60 worth of meals for naval reservists in Maine (link), and $54 spent on levered attachments for holding doors open (link)

[22]The largest medical insurance contract in our dataset is for $20bn (link); the largest non-medical-insurance contract is that for the new F-35 combat aircraft, for which almost $11bn dollars are currently obligated, with the option to increase that to $175bn if needed (link).

and around 950 six-digit NAICS sectors, Table 9). In contrast, the number of recipients is relatively small (around 40,000 per year relative to about 4 million contracts per year).

The constructed monthly spending by transaction/contract is aggregated to county-by-month spending over the period of January 2019 and December 2022 to instrument for county-by-month labor market tightness. The broad geographical coverage of the underlying data ensures broad coverage of counties by the instrument. The distribution of spending across county shows a large degree of variation and a high degree of rightward skewness (Figure 11), the latter is related to the also high degree of skewness in spending by contract.

	25 %	50 %	75 %	95 %	99 %	99.9 %	99.99 %
By contract							
Total obligation (USD)	161	508	2,347	43,700	668,235	13,564,612	140,948,525
Duration (months)	1	1	2	11	30	99	628
By transaction							
Monthly spending (USD)	96	423	1,702	25,272	176,380	1,543,261	10,814,371

Table 7: Cross-sectional summary statistics for procurement contract data

	2019	2020	2021	2022
Outstanding contracts				
Number	3,796,762	3,608,416	3,835,155	3,538,846
Spending USD bns	310	326	352	372
Total obligations USD bns	334	350	345	320
New contracts				
Number	3,362,758	3,188,298	3,454,846	3,168,839
Spending USD bns	79	83	85	87
Total obligations USD bns	159	174	166	149

Table 8: Annual summary statistics for procurement contract data

	2019	2020	2021	2022
Number or recipients	42,670	40,733	38,106	36,739
Number of counties	491	491	489	489
number of states	50	50	50	50
number of sectors	964	959	910	926

Table 9: Coverage of procurement contract data

4.2 Results

Our headline results come from estimating equation (1), instrumenting for county-month labor market tightness using 12 lags of log DoD contract procurement spending. Table 11

Industry	Obligation(bn)	Percent
Aircraft Manufacturing	334.0	21.0
Engineering Services	146.0	9.0
Ship Building and Repairing	136.0	9.0
Guided Missile and Space Vehicle Manufacturing	108.0	7.0
Search, Detection, Navigation, Guidance, Aeronautical, and Nautical System and Instrument Manufacturing	78.0	5.0
Research and Development in the Physical, Engineering, and Life Sciences (except Nanotechnology and Biotechnology)	76.0	5.0
Unknown	73.0	5.0
Other Aircraft Parts and Auxiliary Equipment Manufacturing	71.0	5.0
Commercial and Institutional Building Construction	59.0	4.0
Direct Health and Medical Insurance Carriers	51.0	3.0
Aircraft Engine and Engine Parts Manufacturing	49.0	3.0
Pharmaceutical Preparation Manufacturing	36.0	2.0
Military Armored Vehicle, Tank, and Tank Component Manufacturing	34.0	2.0
Computer Systems Design Services	31.0	2.0
Other Computer Related Services	30.0	2.0
Other Support Activities for Air Transportation	24.0	2.0
Facilities Support Services	22.0	1.0
Biological Product (except Diagnostic) Manufacturing	21.0	1.0
Petroleum Refineries	20.0	1.0
Other Heavy and Civil Engineering Construction	19.0	1.0
Ammunition (except Small Arms) Manufacturing	18.0	1.0
Power Boiler and Heat Exchanger Manufacturing	18.0	1.0
Radio and Television Broadcasting and Wireless Communications Equipment Manufacturing	17.0	1.0
Custom Computer Programming Services	16.0	1.0
All Other Professional, Scientific, and Technical Services	16.0	1.0
Other Guided Missile and Space Vehicle Parts and Auxiliary Equipment Manufacturing	16.0	1.0
Medicinal and Botanical Manufacturing	14.0	1.0
Administrative Management and General Management Consulting Services	13.0	1.0
Highway, Street, and Bridge Construction	12.0	1.0
Service Establishment Equipment and Supplies Merchant Wholesalers	12.0	1.0
Whole sample	1,570.0	100.0

Table 10: Summary statistics: DoD spending, by industry

Table 10 shows top the 30 Department of Defense (DOD) spending by industry in the final dataset. “Obligations” is the total amount of spending promised by the DoD within each industry. ”Percent” express the amount of obligations by industry as a fraction of the total obligations across industry.

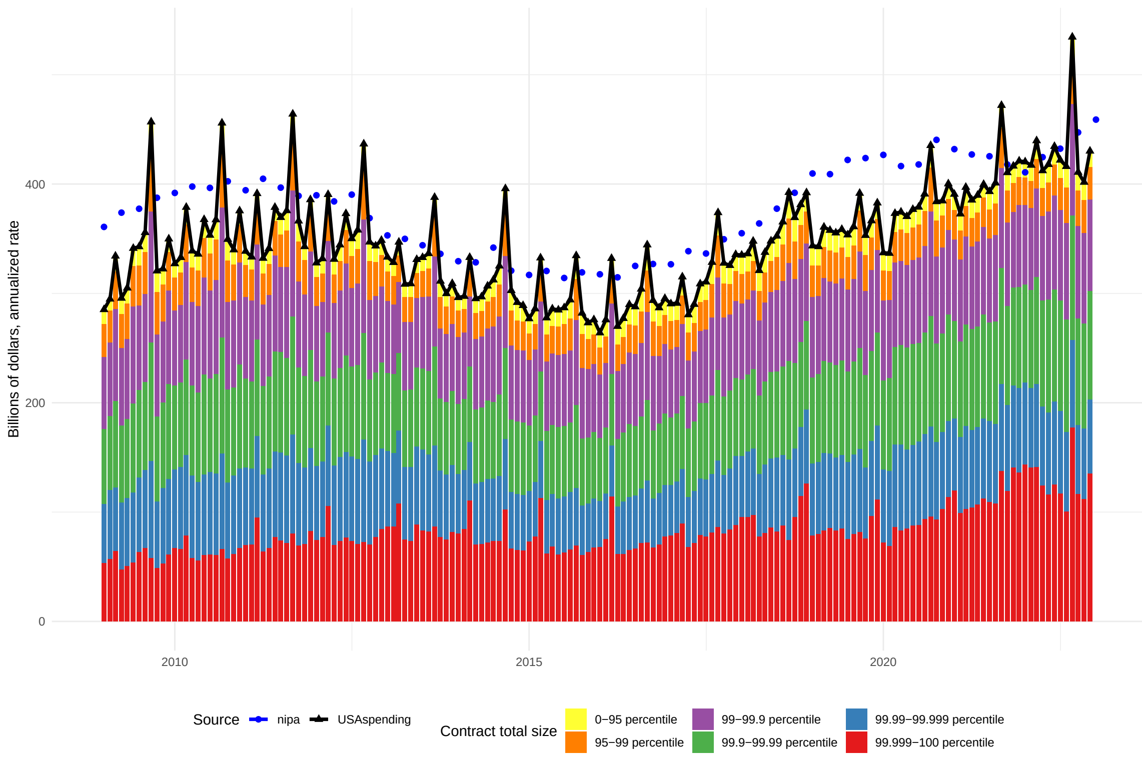

Figure 10: Comparison of NIPA data vs. aggregated from USASpending

Figure compares total spending on Department of Defence procurement from two sources: National Income and Product Accounts (NIPA) and USASpending. The latter is broken down by the total size of the contract, grouped by percentile.

reports summary statistics from the first stage regression. The F statistics are large, all with p-values smaller than 0.001. Instrumenting with fewer lags produces lower F-statistics, indicating that the relationship between changes in federal spending and local labor market tightness may have variable lags. As a check on the direction of the association between the instrument and labor market tightness, Table 11 also reports the sum of the coefficients on the lagged instruments. This is positive, implying that positive local government spending shocks are positively associated with tighter local labor markets, as we should expect. This addresses one of the requirements for an instrument, relevance.

The other requirement for a valid insturment is that it has no direct impact on the outcome variable except through the mechanism proposed. Figure 12 speaks to this point, comparing the DoD expenditure share with the share of workers in the relevant industry. There is almost no overlap. This rules out one way that our instrument could be invalid: if DoD purchases were assigned directly to the industries in which the Homebase firms work, the instrument would not capture a general increase in labor demand, but instead a sector-

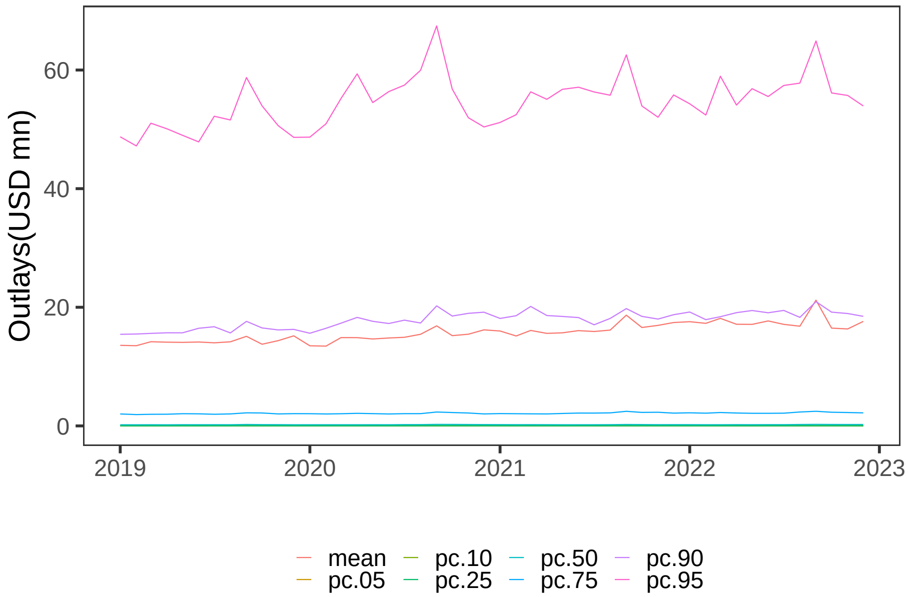

Figure 11: Defense spending by county

Figure plots the distribution of defense spending by county, which is aggregated from the monthly spending at the transaction/contract level constructed from the USASpending.gov dataset.

specific government purchase shock. However, since the firms we study are providing things like catering and retail services, rather than fighter planes and submarines, we can rule this out. Indeed, as Figure 12 makes clear, there is essentially no direct link from our instrument to the firms whose outcomes we study. The only industries with a visible share in Figure 12 are computer programming services, and administrative and management consulting. And so, the DoD procurement shock can be treated as more akin to a general local demand shock than a sector-specific one.

	Full-sample average	Entry-corrected	Double-corrected
Number of observations	19873	19873	11160
Instrument F stat	369	369	276
p value	0.000	0.000	0.000
Sum of coefficients	0.042	0.042	0.022
standard error	(0.024)	(0.024)	(0.028)

Table 11: Summary of instrument first stage regressions

"Instrument F stat" shows the F statistic for the test that all lags of the instrument are jointly zero in the first stage regression. "Sum of coefficients" is the sum of the coefficients on the lagged instruments, and so corresponds to the cumulative impact of a permanent unit increase in the instrument on labor market tightness.

Figure 13 shows the results from the instrumental variables regression. Qualitatively,

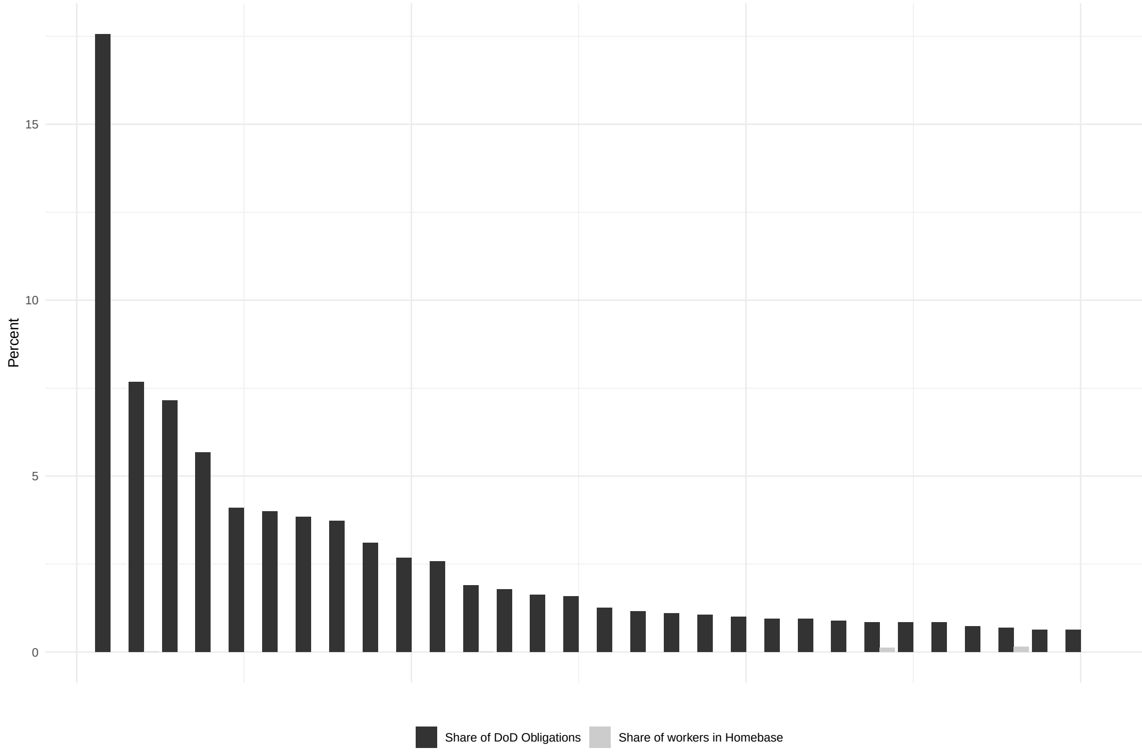

Figure 12: Distribution of DoD spending by industry

Figure plots the spending by the top 30 industries as a share of total DoD spending in our final dataset, which is compared with the same industry's share of workers within our Homebase final dataset.

the results are broadly consistent with OLS: the number of employees, average hours, firm survival rates and employee match rates all fall. Wages rise, although statistical significance for wages remains elusive.

Beyond this overall agreement, the IV results add more detail to the firm responses to local labor market shocks. Most notably, the shock to labor markets themselves has a hump-shaped response, suggesting that increases in labor market tightness induced by local government spending shocks are overall more persistent than the average increase labor market tightness. Nevertheless, this does not necessarily mean that we identify a fundamentally different shock with a different character and transmission mechanism. To assess whether this is true, Figure 18 in Appendix B computes the impulse responses for the instrumental variables, rescaling them so that the cumulative labor market impulse equals that from the OLS estimates. In this case, the difference in the other responses goes away almost entirely. This suggests that the shock identified by instrumental variables is essentially the same as that in ordinary least squares, with the same transmission, just that the magnitude of the shocks identified using the instrument is a little larger.

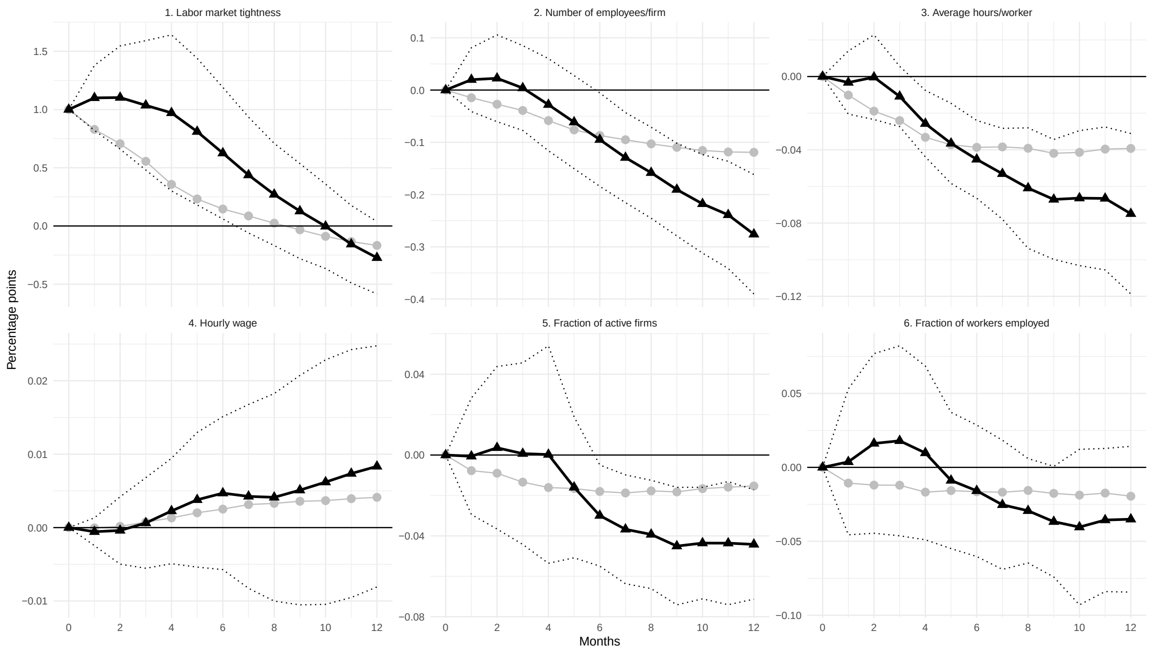

Figure 13: Headline results: Instrumental variables

Panel shows responses of outcome variables to a one percent increase in labor market tightness. Dotted lines show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. Instrument is 12 lags of local Department of Defense procurement spending. Grey circles are the OLS results

4.3 Robustness

In Appendix B we show that the instrumental variables estimates are robust to a the same factors as our OLS estimates, including specification, firm size, and other sample restrictions.

5 Conclusion

In this paper, we use proprietary micro data from Homebase—a proprietary dataset covering nearly 9 million workers at 1 million firms—to present substantial empirical evidence that small firms are negatively affected by tight labor market conditions, consistent with the "job-ladder" theory. We find that firms in our sample reduce employment and hours per worker while increase wages following an increase in local labor market tightness. These effects persist for at least a year and are statistically significant. The number of firms and worker-firm matches observed also decreases, consistent with higher firm failure rates, although this could be due to other factors.

A key methodological contribution of this paper is to correct for composition bias. Without the correction, the impulse responses would paint a much rosier picture for small firms, suggesting that they barely reduce hours and can increase wages by almost twice as much. Controlling for dynamic composition bias is also important in the context of thinking about the slope of the wage Phillips curve. Our results on the impact of labor market tightness on wage growth is not consistent with predictions from traditional labor market models where wages are determined by the bargaining power of job stayer. Instead, our results suggest that wage growth is driven by new entrants, consistent with a job-ladder model.

We confirm the interpretation of our results, namely that the outcomes we observe are due to local labor market shocks, using an instrumental variable approach. This involves using detailed USASpending data on defense contract awards, making this paper one among very few studies that have done so. The instrumental variable results are very similar to the baseline.

A fruitful venue for future empirical research is to compare the labor market responses of large and small firms and examine whether firm size fully accounts for the differential responses, or if other factors remain. From a conceptual viewpoint, the empirical results would benefit from a theoretical framework to further explain some of the patterns we uncover, namely that small firms have a tough time in tight labor markets, building on work already done in this area (e.g., Moscarini and Postel-Vinay (2023)).

References

Altonji, J., Z. Contractor, L. Finamor, R. Haygood, I. Lindenlaub, C. Meghir, C. O'Dea, D. Scott, L. Wang, and E. Washington, "Employment Effects of Unemployment Insurance Generosity During the Pandemic," *Working Paper, Tobin Center for Economic Policy, Yale University*, 2020.

Auerbach, A., Y. Gorodnichenko, and D. Murphy, "Local Fiscal Multipliers and Fiscal Spillovers in the USA," *IMF Economic Review*, 2020, *68* (1), 195–229.

Auerbach, Alan, Yuriy Gorodnichenko, and Daniel Murphy, "Local fiscal multipliers and fiscal spillovers in the USA," *IMF Economic Review*, 2020, *68*, 195–229.

Autor, D., A. Dube, and A. McGrew, "The Unexpected Compression: Competition at Work in the Low Wage Labor Market," *NBER Working Paper*, 2023, (31010).

Bartik, W.A., M. Bertrand, F. Li, and J. Rothstein, "Measuring the Labor Market at the Onset of the COVID-19 Crisis," *NBER Working Paper*, 2020, (27613).

Blanchard, Olivier, Alex Domash, and Lawrence Summers, "Bad News for the Fed from the Beveridge Space," *Petreson Institute for International Economics Policy Brief*, 2022, *22* (7).

_ **and Jordi Galí**, "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment," *American Economic Journal: Macroeconomics*, 2010, *2*, 1–30.

_ **and Peter A. Diamond**, "The Beveridge Curve," *Working Paper*, 1989, *1*, 1–76.

CEA, "Investing in America Means Investing in America's Small Businesses," Technical Report 2023.

Chen, Sophia and Do Lee, "The Post-Pandemic Paycheck: Wage and Earning Dynamics from Micro Data," *IMF Working Paper*, 2024.

Chetty, R., J. Friedman, N. Hendren, and M. Stepner, "The Economic Impacts of COVID-19: Evidence from a New Public Database Built using Private Sector Data," *Quarterly Journal of Economics*, 2023, *forthcoming.*

Christiano, Larence J, Martin S Eichenbaum, and Mathias Trabandt, "Unemployment and Business Cycles," *Econometrica*, 2016, *84* (4), 1523–1569.

Cox, Lydia, Gernot Müller, Ernesto Pasten, Raphael Schoenle, and Michael Weber, "Big G," Technical Report, National Bureau of Economic Research 2020.

Crouzet, N. and N.R. Mehrotra, "Small and Large Firms over the Business Cycle," *American Economic Review*, 2020, *110* (11), 3649–3601.

Duval, Mr Romain A, Yi Ji, Longji Li, Myrto Oikonomou, Carlo Pizzinelli, Mr Ippei Shibata, Alessandra Sozzi, and Marina M Tavares, *Labor market tightness in advanced economies*, International Monetary Fund, 2022.

Dvorkin, M. A. and M. Isaacson, "Tracking Wage Inflation in Real Time," *Federal Reserve Bank of St Louis, On the Economy Blog*, 2022.

Galí, Jordi, Frank Smets, and Rafael Wouters, "Unemployment in an Estimated New Keynesian Model," *NBER Macroeconomics Annual*, 2011, *26* (1), 329–360.

Gertler, M. and S. Gilchrist, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics*, 1994, *109* (2), 309–40.

Granja, J., C. Makridis, C. Yannelis, and E. Zwick, "Did the paycheck protection program hit the target," *Journal of Financial Economics*, 2022, *145* (3), 725–761.

Hooker, Mark A and Michael M Knetter, "The effects of military spending on economic activity: Evidence from state procurement spending," *Journal of Money, Credit, and Banking*, 1997, pp. 400–421.

Jordà, Òscar, "Estimation and inference of impulse responses by local projections," *American economic review*, 2005, *95* (1), 161–182.

Karahan, F., R. Michaels, B. Pugsley, A. Şahin, and R. Schuh, "Do Job-to-Job Transitions Drive Wage Fluctuations over the Business Cycle," *American Economic Review Papers and Proceedings*, 2017, *107* (5), 353–57.

Kurmann, A., E. Lale, and L. Ta, "Measuring Small Business Dynamics and Employment with Private-Sector Real-Time Data," *Working Paper*, 2023.

Kydland, Finn E and Edward C Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica*, 1982, *50*, 1345–1370.

Morales-Jiménez, Camilo, "The Cyclical Behavior of Unemployment and Wages under Information Frictions," *American Economic Journal: Macroeconomics*, 2021, *14* (1), 301–331.

Mortensen, Dale T, "Markets with Search Friction and the DMP Model," *The American Economic Review*, 2011, *101* (4), 1073–1091.

_ **and Christopher A Pissarides**, "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, 1994, *61* (3), 397–415.

Moscarini, G. and F. Postel-Vinay, "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment," *American Economic Review*, 2012, *102* (6), 2509–39.

_ **and** _ , *Journal of Labor Economics*, 2016, *34* (Did the Job Ladder Fail After the Great Recession), 55–93.

_ **and** _ , "Wage Posting and Business Cycles," *American Economic Review Papers and Proceedings*, 2016, *106* (5), 208–13.

— **and** — , "The Relative Power of Employment-to-Employment Reallocation and Unemployment Exits in Predicting Wage Growth," *American Economic Review Papers and Proceedings*, 2017, *107* (5), 364–68.

— **and** — , "The Job Ladder: Inflation vs. Reallocation," *NBER Working Paper*, 2023, (31466).

Pissarides, Christopher A, *Equilibrium unemployment theory*, MIT press, 2000.

— , "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer," *Econometrica*, 2009, *77* (5), 1339–1369.

Ramey, Valerie A, "Can government purchases stimulate the economy?," *Journal of Economic Literature*, 2011, *49* (3), 673–685.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *The American Economic Review*, 2005, *95* (1), 25–49.

A Further robustness checks

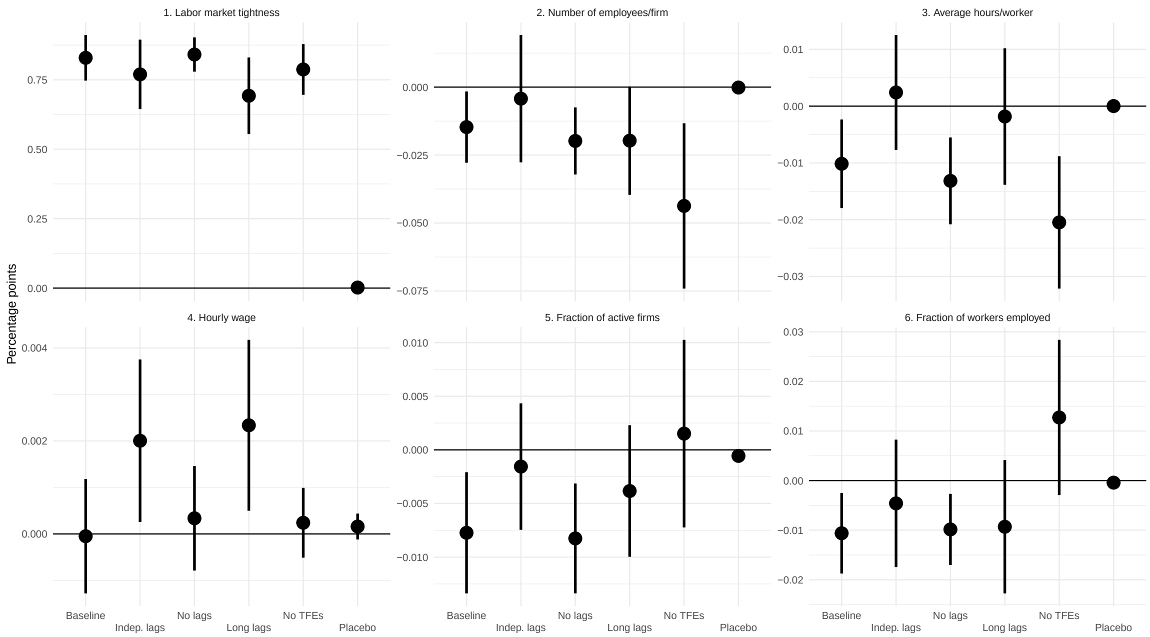

Figure 14: Robustness of OLS estimates: alternate specifications, $h = 1$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 1$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Indep. lags* adds 3 lags of the independent variable, *Long lags* includes 12 lags of both the independent and dependent variables, *No lags* omits lags of both dependent and independent variables, *No TFEs* omits time fixed effects, and *Placebo* replaces the independent variable with a white noise proces withh the same standard deviation as the uneployment-vacancy ratio in the data.

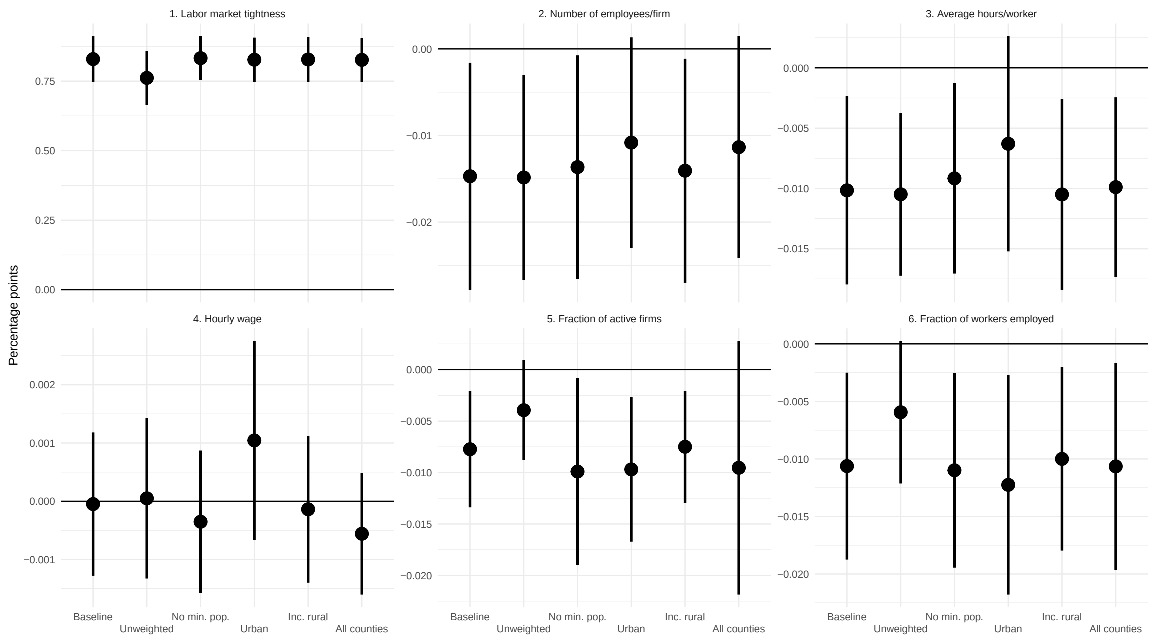

Figure 15: Robustness of OLS estimates: alternate specifications, $h = 1$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 1$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Unweighted* removes population weighting, *No min. pop.* removes the minimum county population requirement, *Urban* uses only counties with a CDC Urban score of 1 or 2, *Inc. rural* includes counties with CDC urban score up to 6 (i.e. no restriction), and *All counties* includes all counties without restriction on population or urban/rural character.

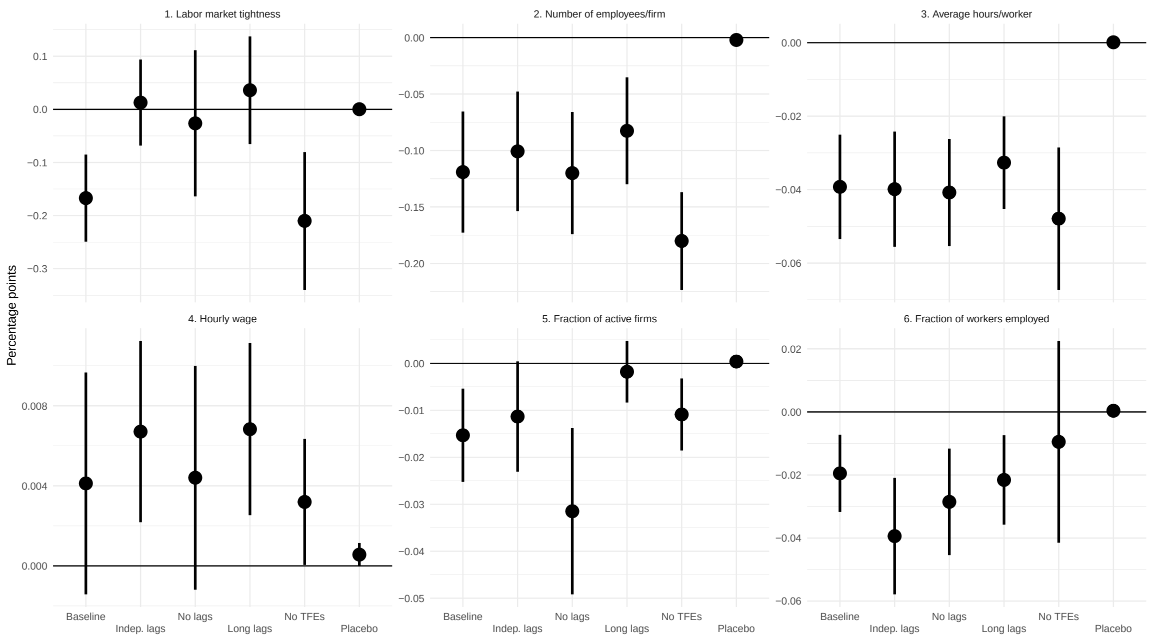

Figure 16: Robustness of OLS estimates: alternate specifications, $h = 12$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 12$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Indep. lags* adds 3 lags of the independent variable, *Long lags* includes 12 lags of both the independent and dependent variables, *No lags* omits lags of both dependent and independent variables, *No TFEs* omits time fixed effects, and *Placebo* replaces the independent variable with a white noise proces withh the same standard deviation as the uneployment-vacancy ratio in the data.

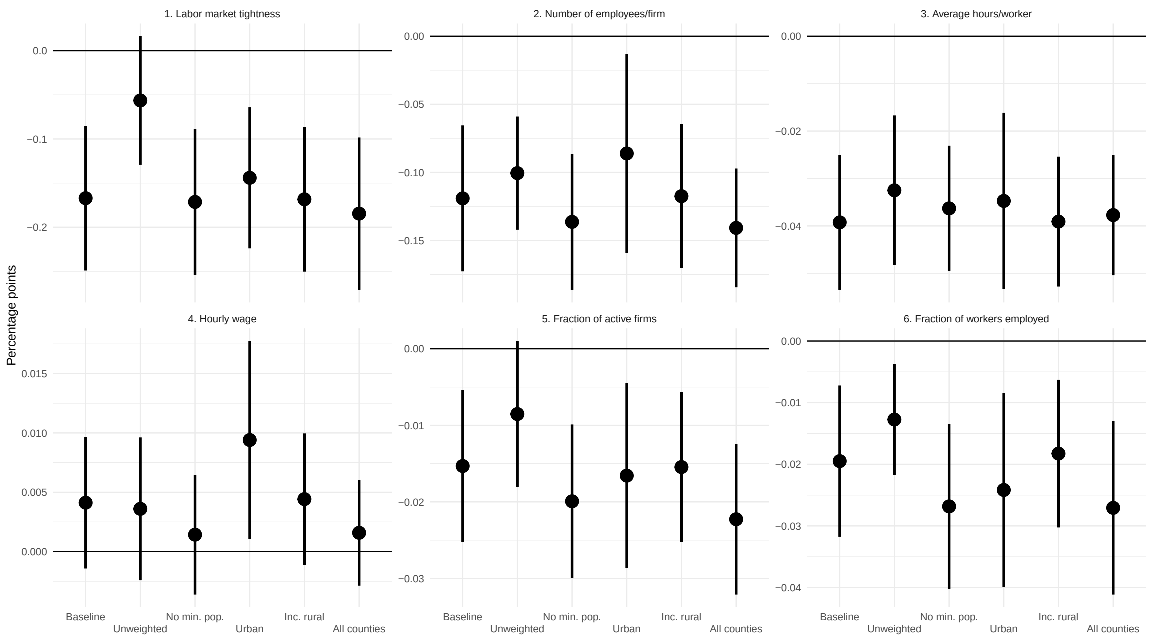

Figure 17: Robustness of OLS estimates: alternate specifications, $h = 12$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 12$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Unweighted* removes population weighting, *No min. pop.* removes the minimum county population requirement, *Urban* uses only counties with a CDC Urban score of 1 or 2, *Inc. rural* includes counties with CDC urban score up to 6 (i.e. no restriction), and *All counties* includes all counties without restriction on population or urban/rural character.

B Instrumental Variables

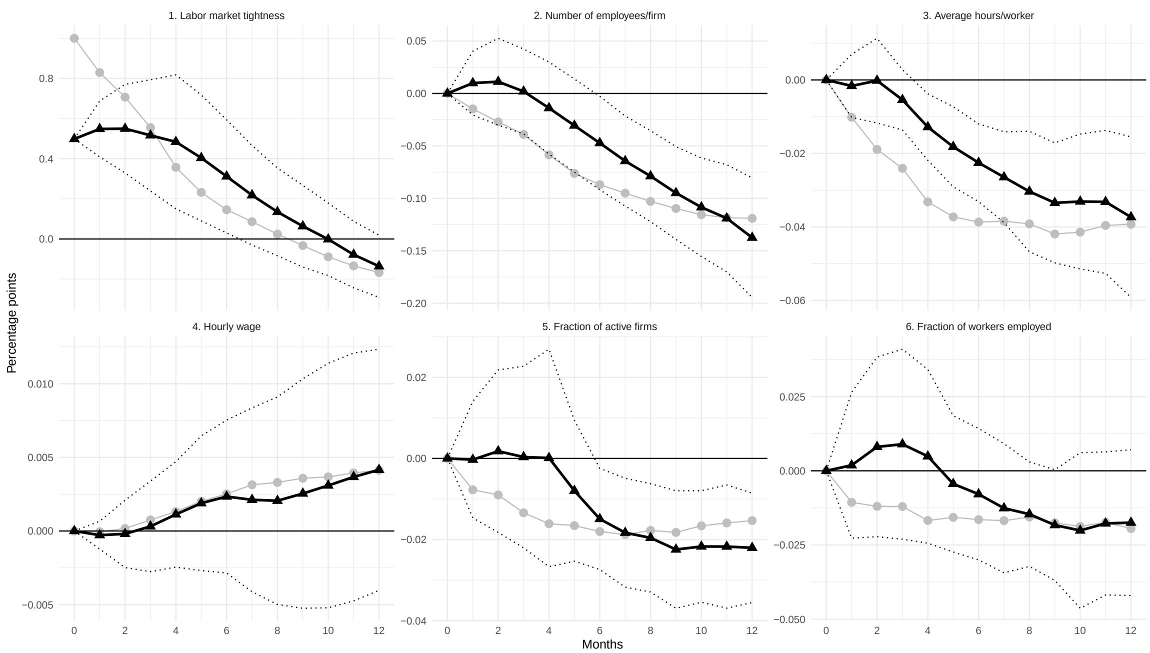

Figure 18: Headline results: Instrumental variables, rescaled

Panel shows responses of outcome variables to a one percent increase in labor market tightness. Dotted lines show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. Instrument is 12 lags of local Department of Defense procurement spending. Grey circles are the OLS results. Instrumental variables estimates are rescaled by a constant so that the the cumulative labor market tightness responses are identical.

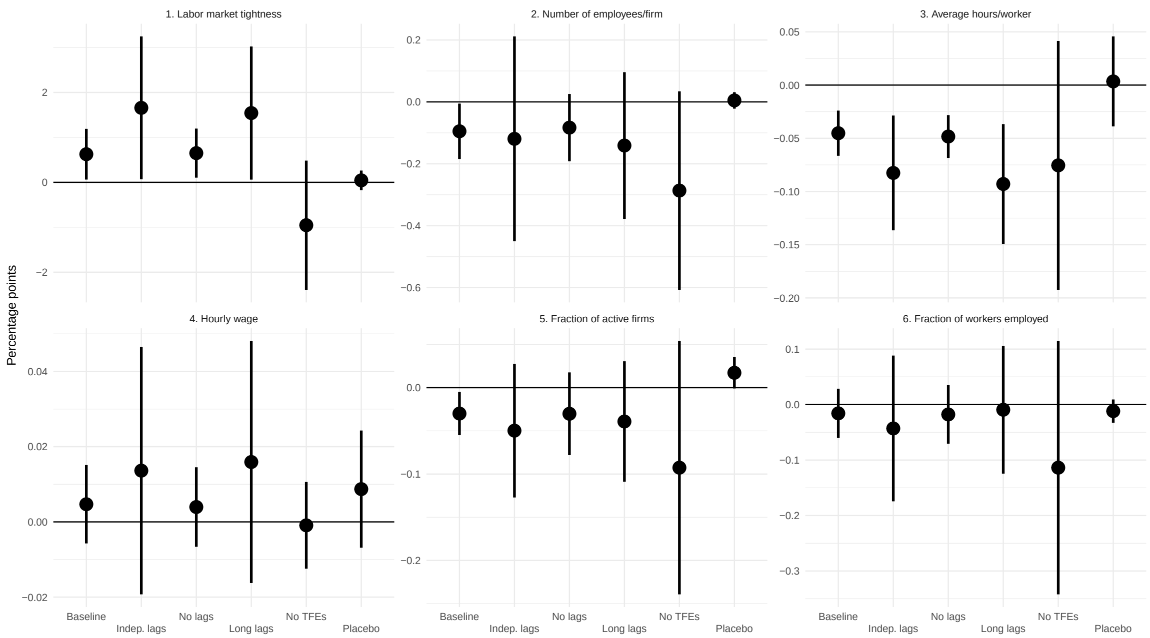

Figure 19: Robustness of IV estimates: alternate specifications, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Indep. lags* adds 3 lags of the independent variable, *Long lags* includes 12 lags of both the independent and dependent variables, *No lags* omits lags of both dependent and independent variables, *No TFEs* omits time fixed effects, and *Placebo* replaces the independent variable with a white noise process with the same standard deviation as the unemployment-vacancy ratio in the data.

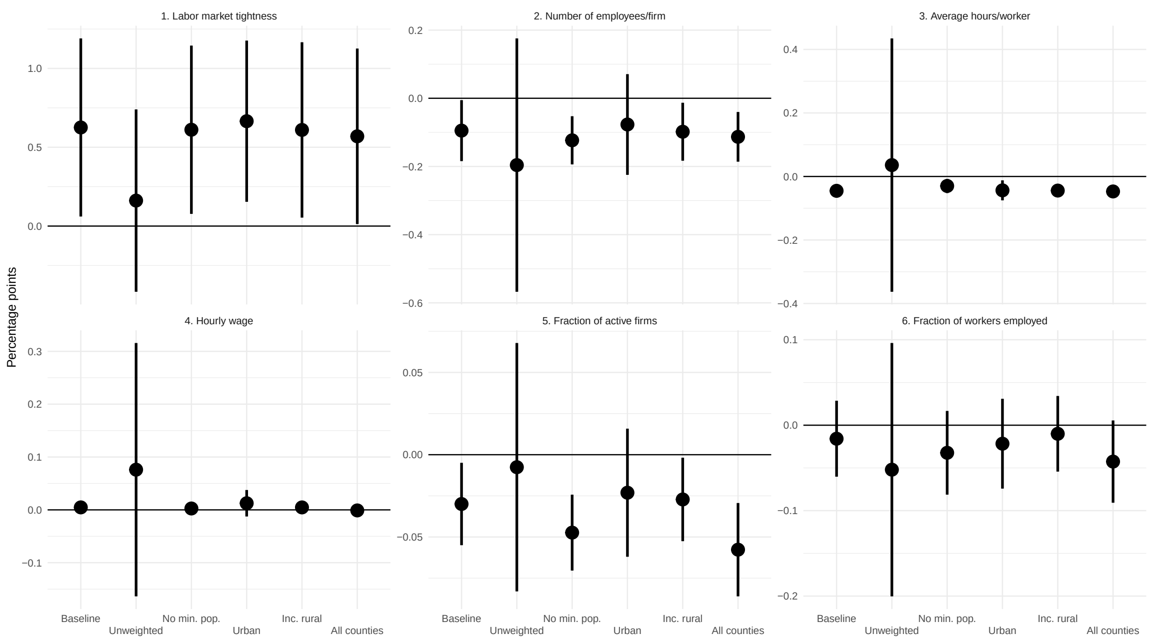

Figure 20: Robustness of IV estimates: alternate sample restrictions, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. *Unweighted* removes population weighting, *No min. pop.* removes the minimum county population requirement, *Urban* uses only counties with a CDC Urban score of 1 or 2, *Inc. rural* includes counties with CDC urban score up to 6 (i.e. no restriction), and *All counties* includes all counties without restriction on population or urban/rural character.

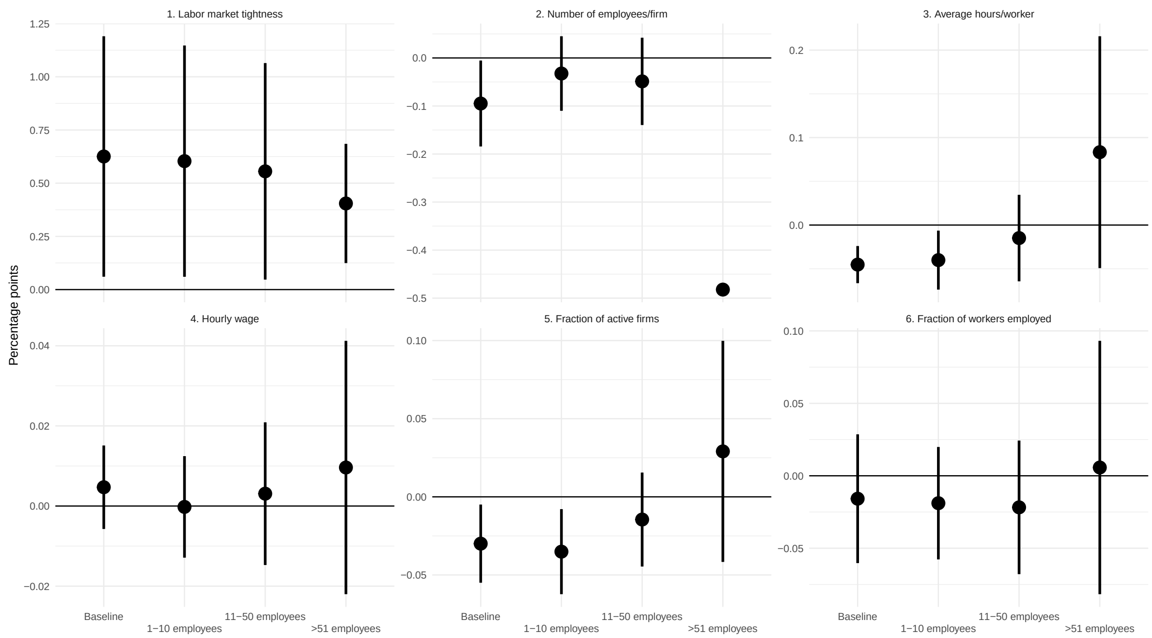

Figure 21: Robustness of IV estimates: firms size, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Observations are weighted by county-level populations. Estimates vary by firm size categories, as indicated in the chart.

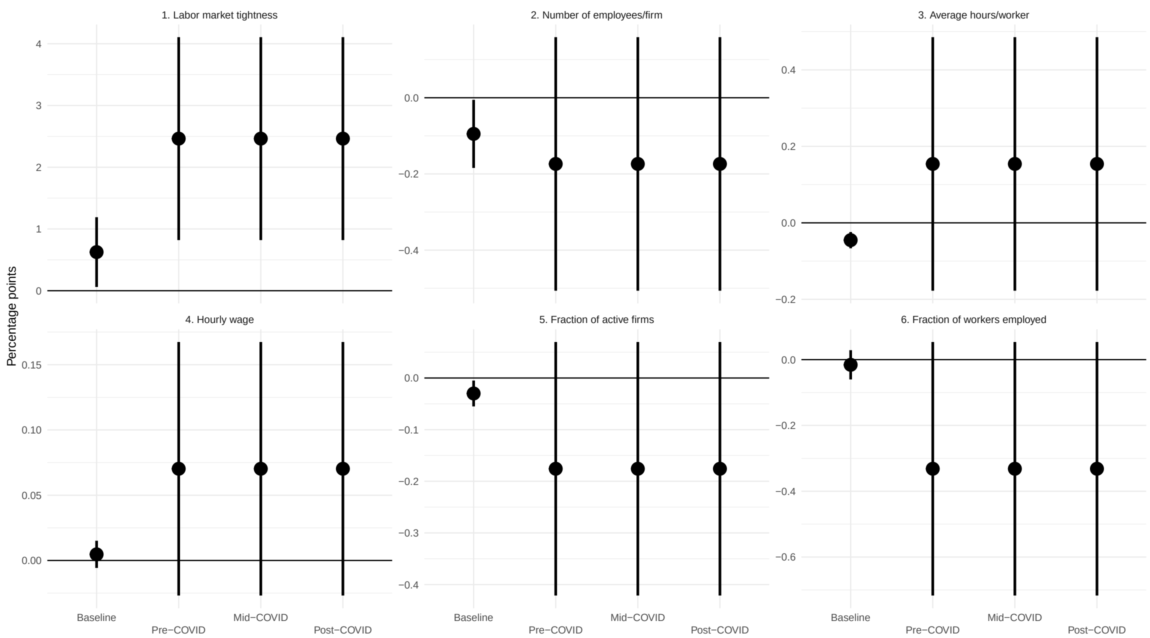

Figure 22: Robustness of IV estimates: alternate sample periods, $h = 6$

Panel shows responses of outcome variables to a one percent increase in labor market tightness at horizon $h = 6$ only. Ranges show 95 percent confidence intervals computed using robust standard errors double-clustered at the county and month. Country-month outcomes are calculated from Homebase microdata using dynamic sample double-correction in panels 2-4 and entry-correction panel 5 and 6. Estimates vary by period, *Pre-COVID* is Jan 2019-Feb 2020, *Mid-COVID* is Mar 2020-Jun 2021, and *Post-COVID* is Jul 2021-Nov 2022.

PUBLICATIONS

The Dynamic Effects of Local Labor Market Shocks on Small Firms in The United States

Working Paper No. WP/2024/063