

Valuing Labor Market Power: The Role of Productivity Advantages

Bryan Seegmiller*

First Version: November 2021

Current Version: September 2025

Abstract

Using idiosyncratic stock returns, I estimate heterogeneous firm-level labor supply elasticities by labor productivity, worker skill, and time. After accounting for the mitigating impact of adjustment costs, I use these elasticity estimates to quantify how wage markdowns affect the following: a wide cross-sectional labor share spread by productivity; the public firm aggregate labor share decline from 1991-2014; and productive firms' high profits and valuations, despite low investment. Overall, profits from wage markdowns are worth 20-25% (4%) of aggregate capital income (revenues). Productive firms' labor market power over skilled workers plays a central role in these patterns.

*Kellogg School of Management. I thank my advisors Leonid Kogan, David Autor, Larry Schmidt, and Dimitris Papanikolaou for extremely helpful guidance and continued support. I also thank Daron Acemoglu, Hui Chen, Taha Choukhmane, Maryam Farboodi, Jack Favilukis (discussant), Dan Greenwald, Simon Jäger, Eben Lazarus, Debbie Lucas, Chris Palmer, Jonathan Parker, Antoinette Schoar, Anna Stansbury, David Thesmar, Adrien Verdelhan, Jessie Wang (discussant), Yufeng Wu (discussant), Emil Verner, Haoxiang Zhu, and Luigi Zingales, as well as participants at numerous seminars and conferences for helpful comments. The U.S. Census Bureau reviewed this data product for unauthorized disclosure of confidential information and approved the disclosure avoidance practices applied to this release (approval numbers CBDRB-FY21-163, CBDRB-FY21-POP001-0178, CBDRB-FY21-POP001-0232, CBDRB-FY22-SEHSD003-004, CBDRB-FY25-0268, CBDRB-FY25-0418).

Market power shifts the value of what a firm produces away from its workers and toward the owners of its capital. Much of the literature studying market power’s role in recent macroeconomic trends—the large rise in the value of the aggregate stock market and in firm profitability (Greenwald, Lettau, and Ludvigson, 2021; Barkai, 2020); high firm valuations alongside low rates of investment (Gutierrez and Philippon, 2017); and the decline in the share of output accruing to labor (Karabarbounis and Neiman, 2013)—has focused on product markets (De Loecker, Eeckhout, and Unger, 2020; Autor, Dorn, Katz, Patterson, and Van Reenen, 2020). Similar to Berger, Herkenhoff, and Mongey (2022), this paper shifts the focus to labor markets, showing that firms with high labor productivity enjoy a competitive advantage in the form of substantial wage markdowns from marginal revenue product. Since labor compensation is the largest production cost for most firms, paying workers below their marginal revenue product creates significant value for capital owners. Holding equilibrium quantities fixed, I estimate that wage markdowns are worth about a fifth to a quarter of total capital income to the publicly traded U.S. corporations in my sample, and an even higher share for the most productive firms. This “superstar firms” view of labor market power helps explain cross-sectional differences in firm-level profitability, labor shares, and valuation ratios, as well as time-series changes in the aggregate labor share.

Getting to these numbers requires estimating the slope of the labor supply curve facing individual firms, which in turn requires shocks that move a firm’s desired labor demand without moving the labor supply curve it faces. My starting point, developed in section 2, is that *any* such shock contains valid identifying variation: pure TFP shocks, product demand shocks, technological change, and labor cost shifters that don’t operate through the labor supply curve all qualify. Since all of these affect stock returns, idiosyncratic stock returns are a potentially powerful labor demand shifter. Returns impound a wide range of news about a firm’s prospects, so they reflect far more of the variation that moves labor demand than any single observable shock does, and they are available for every public firm in every year. While stock returns have been extensively studied in various other contexts, this paper is the first to systematically examine the relationship between a firm’s own stock returns and its labor demand.¹ In contrast to Berger et al. (2022), who capture heterogeneity in firms’ labor market power through the lens of a structural model of local labor market competition, my approach is to estimate supply elasticity heterogeneity directly and along several dimensions at once: across time, across firm types, and across worker skill levels.

I bring this framework to employer-employee matched data from the Longitudinal Employer-Household Dynamics database (LEHD), linked to publicly traded firms in Compustat, over 1991-2014.² This sample includes both wide variation in labor productivity and

¹While Kogan, Papanikolaou, Schmidt, and Song (2020) relate wage growth to stock returns, their focus is on technological innovation and income risk.

²The sample period is set by data availability. LEHD coverage is thin at the start of the 1990s and becomes comprehensive by the mid- to late 1990s, and the LEHD extract available for this project ends in 2015; the

many of the world’s most productive firms. Consistent with the view that stock returns contain unexpected information about labor demand, I show, using a series of local projection specifications, that firm-specific returns generate immediate, persistent, and strongly positive effects on forward growth in average wages, employment, and productivity, while showing essentially no relation to past growth.

The remaining identification question is whether labor supply shocks account for enough of the variation in stock returns to materially affect the estimated elasticity. Section 2 shows why market-level shocks are a particularly important potential source of supply variation, while section 3.4 demonstrates that estimates are quantitatively stable across a wide range of controls for common shocks and alternative definitions of labor market boundaries. Estimates also change little when I control for observable firm-specific supply and job-characteristic changes—including unionization, non-compete enforceability, and hours variation—or when I replace own-firm returns with three alternative labor demand instruments. Taken together, these exercises indicate that any supply-side contamination of stock returns is quantitatively minor. Finally, on top of their limited empirical importance in practice, I show that as long as firms do not typically cut amenities sharply when their own productivity improves, any remaining firm-specific amenity shocks push the estimated elasticity upward, making my estimates of labor market power a lower bound; internet appendix IA.3 derives the condition.

Using the ratio of the changes in one-year growth rates in employment and wages induced by stock returns, I find an implied labor supply elasticity of about 2.5 for my pooled, full-sample estimate. Two dimensions of heterogeneity in this elasticity drive everything that follows. The first is worker skill. Leveraging the worker-level resolution in the LEHD, I proxy for worker skill using the size of the worker-specific fixed effect in an Abowd, Kramarz, and Margolis (1999) style wage decomposition, and I find that skilled workers exhibit far lower labor supply elasticities than low-skill workers: workers in the bottom 40% of the skill distribution have a labor supply elasticity of around 8, while workers in the top 20% of the skill distribution have a supply elasticity of 1.2, implying greater market power over skilled workers. The second dimension is firm labor productivity. In specifications that impose homogeneous worker types, firms in the top productivity quartile have a labor supply elasticity of 1.2, while firms in the bottom labor productivity quartile have an elasticity of 4.3. Interacting the two, elasticities decline across firm productivity rank regardless of worker skill, so that highly-skilled workers employed at highly-productive firms are the least elastic group of all.

I sort on cross-sectional labor productivity differences for two reasons. First, heterogeneous market power is a classic source of marginal revenue product dispersion (Hsieh and Klenow, 2009; Peters, 2020), making this a natural dimension to test for differences in labor market

one-year growth rates my specification requires yield annual elasticity estimates for 1992-2013. Section 1 describes the data and the sample construction in detail.

power. Second, stylized facts for firms with higher labor productivity strongly suggest that they do earn high economic rents relative to less productive firms: they have low labor shares, high profitability, and high valuation ratios, but they do not have high investment rates. Since a lower firm-specific elasticity generates greater profits from market power, productive firms' low elasticities confer a labor market competitive advantage. Correspondingly, wage markdowns implied by elasticity estimates across the productivity distribution are strongly associated with these firm outcomes.

My estimates also track important changes in the competitive landscape of the US economy over time. Supply elasticities are uniformly lower across the productivity distribution in the second half of my sample period (2003-2014) than in the first half (1991-2002), which, viewed through the lens of a wage-posting monopsonistic model, implies a secular increase in labor market power. Applying my method at higher frequency using 3-year moving-window average estimates, the cross-sectional gap in elasticity-implied wage markdowns between the most- and least-productive firms widens over time, and this gap closely tracks a widening spread in the average log labor share, which suggests these estimates reflect genuine differences in market power in the labor market. Measures of local labor market concentration, by contrast, have trended downward over this period and hence cannot explain these aggregate trends.³

Translating elasticities into markdowns takes one further step. In section 4 I show that static labor supply elasticity estimates are highly informative, but still insufficient, for recovering firms' actual wage markdowns when labor determination is dynamic due to adjustment costs. Prior research has accumulated evidence that labor is costly to adjust (Kline, Petkova, Williams, and Zidar, 2019; Jäger, Heining, and Lazarus, 2024; Belo, Lin, and Bazdresch, 2014). Labor adjustment costs result in an additional term in the firm's Euler equation that tends to mitigate monopsony power relative to a static setting, so that the wage markdown decomposes into the typical standard markdown plus an additional dynamic component. Because the dynamic component tends to mitigate monopsony power, ignoring these forces can cause one to make erroneously large inferences about the size of profits firms earn from market power. Importantly, this Euler equation can be estimated directly from the data: a simple function of expected firm hiring rates captures expectations about future adjustment costs.

Section 5 carries out that estimation and delivers the paper's quantitative results. Estimating the framework with homogeneous worker types, profits from wage markdowns represent about 25% of aggregate capital income on average, 28% among the more productive half of firms and 13% among the less productive half; allowing for skill heterogeneity in markdowns gives 23% of capital income (22% and 20% respectively against intangible-adjusted capital income). The dynamic correction operates almost entirely through skilled workers. Cali-

³Berger et al. (2022), Rinz (2022), and Lipsius (2018) all find declining local labor market concentration over the past several decades using Census administrative data.

brated adjustment costs imply that low- and middle-skilled workers’ wage markdowns are only slightly adjusted for dynamics, but the dynamic adjustments are of first-order quantitative importance for skilled workers’ wage markdowns: while skilled workers have an average supply elasticity of around 1.2, firms’ average skilled-worker wage markdowns are consistent with a supply elasticity that’s above 4. Even after this adjustment, skilled workers face greater wage markdowns in general, and represent nearly 70% of profits generated from labor market power. Labor adjustment costs also increase in the time series, which dampens the effect that declining supply elasticities have on the aggregate labor share; increasing labor market power still explains about a third of the public firm labor share decline in my homogeneous specification, and three-fifths with skill heterogeneity. In the cross-section, differing wage markdowns explain about half of a wide gap in average log labor shares between productive and unproductive firms (a third with skill heterogeneity), and a similar share of gaps in firms’ operating profitability as measured by return on assets. Feeding counterfactual profit rates without labor market power through a [Peters and Taylor \(2017\)](#) intangible-adjusted Tobin’s Q decomposition, while holding firm investment rates constant, closes much of a cross-sectional “investment gap” ([Gutierrez and Philippon, 2017](#); [Crouzet and Eberly, 2021](#)) between productive and unproductive firms.

The through-line of these results is the skilled labor market. Skilled workers are the least elastic group in the data; they generate the large majority of the value firms capture from labor market power; and they are the only group for whom dynamic adjustment substantially changes the picture, which suggests that adjustment in the skilled worker labor market is difficult both for these workers and for the firms that employ them. My estimates do not identify the exact frictions responsible, though firm-specific human capital of skilled workers may play a central role. Such a mechanism would also help explain why concentration measured over narrow geographic boundaries misses so much of what these estimates pick up.

The rest of the paper proceeds as follows. Section 1 describes the data and sample. Section 2 develops the estimation framework and the case for stock returns as a labor demand shifter. Section 3 presents the supply elasticity estimates and robustness checks. Section 4 derives wage markdowns under dynamic monopsony, and section 5 quantifies the implications for firm cash flows, labor shares, and valuations. Section 6 concludes.

Related Literature

This paper adds to a growing literature in macro-finance which examines recent macroeconomic trends in competition, firm valuations/operating performance, and labor shares.⁴ Most related to my paper is [Hartman-Glaser, Lustig, and Xiaolan \(2019\)](#), who document that large firms have obtained an increasing share of their sales as capital income over the

⁴These include [Barkai \(2020\)](#), [Grullon, Larkin, and Michaely \(2019\)](#), [Covarrubias, Gutiérrez, and Philippon \(2020\)](#), [Greenwald et al. \(2021\)](#), [Corhay, Kung, and Schmid \(2020\)](#), and [Farhi and Gourio \(2018\)](#).

past several decades and propose an explanation based on firms providing insurance to their workers against increased risk. I similarly find that labor shares have decreased by more among highly productive firms, but I differ by both directly linking this to productive firms marking down wages further from marginal product, and quantifying the implications for profitability and valuations. A parallel area of research in macroeconomics connects the secular decline in labor shares with changes in market competition. I differ from most of this literature in primarily focusing on the role of labor market power instead of product market power.⁵ My findings imply that increases in the labor market power of the most productive firms have also played a crucial role in my sample of publicly-traded companies.

This paper adds to a growing body of work in financial economics that examines the interplay between labor markets and firm financial outcomes and decision making.⁶ My findings suggest that accounting for imperfect competition in labor markets may be of first-order importance for future research in this area.

Finally, this paper builds on recent research in labor economics. A number of prior papers estimate the elasticity of supply to the firm to infer the extent of monopsony power;⁷ another related empirical literature has examined how labor market concentration affects wages.⁸ My paper differs from others in this literature by quantifying the impact that labor adjustment costs have on wage markdowns. I also estimate heterogeneity in labor supply elasticities based on firm characteristics, focusing on labor productivity in particular. While prior work has estimated how supply elasticities that vary by time or by worker skill type (see, for instance Webber, 2022; Bachmann, Demir, and Frings, 2022), to my knowledge, my paper is the first to estimate heterogeneous supply elasticities by worker skill, interacted with both firm productivity and time. Finally, a large literature estimates the passthrough of firm-specific shocks to worker outcomes.⁹ These papers typically analyze the response of worker pay to firm-specific shocks, whereas my focus is on estimating the ratio of the firm-level employment and wage responses to stock return shocks in order to estimate the elasticity of the supply

⁵See Autor et al. (2020), De Loecker et al. (2020), Covarrubias et al. (2020) and Kehrig and Vincent (2021) for examples. One exception is Stansbury and Summers (2020), who argue that worker bargaining power has played an important role.

⁶References in asset pricing include Eisfeldt and Papanikolaou (2013), Belo et al. (2014), Donangelo, Gourio, Kehrig, and Palacios (2019), Kuehn, Simutin, and Wang (2017), Liu (2019), Favilukis and Lin (2015), Favilukis, Lin, and Zhao (2020); see Matsa (2010), Jeffers (2019), Mueller, Ouimet, and Simintzi (2017), Kim (2020), Shen (2021), and Ferres, Kankanalli, and Muthukrishnan (2023) for examples in corporate finance.

⁷See Lamadon, Mogstad, and Setzler (2022), Berger et al. (2022), Kroft, Luo, Mogstad, and Setzler (2025), Bassier, Dube, and Naidu (2022), and Ransom and Sims (2010) for a few examples; Card (2022) surveys this body of work. Alternatively, Yeh, Macaluso, and Hershbein (2022) and Mertens (2023) estimate labor market power through a production function estimation approach, with the latter finding that “superstar” German manufacturing firms have enjoyed increased market power in recent decades.

⁸See Rinz (2022), Benmelech, Bergman, and Kim (2022), Schubert, Stansbury, and Taska (2024), and Jarosch, Nimczik, and Sorkin (2024), Azar, Marinescu, and Steinbaum (2022) for example.

⁹See Abowd and Lemieux (1993), Van Reenen (1996), Guiso, Pistaferri, and Schivardi (2005), Kline et al. (2019), Kogan et al. (2020), Chan, Salgado, and Xu (2021), Garin and Silverio (2020), Friedrich, Laun, Meghir, and Pistaferri (2019), Balke and Lamadon (2020), Howell and Brown (2020).

curve that the firm faces.

In a related prior paper, [Gouin-Bonenfant \(2022\)](#) argues theoretically for labor productivity dispersion as a determinant of labor market power in a search and matching model, and he shows that an increase in productivity dispersion can match key aggregate empirical moments, including a decline in the aggregate labor share. I differ from [Gouin-Bonenfant \(2022\)](#) in providing direct empirical estimates of firm-level monopsony power as implied by labor supply elasticities. I also differ in my focus on quantifying the value of the capital income that is derived from wage markdowns, and the role of skilled workers in driving it. Thus my findings complement and add to those of [Gouin-Bonenfant \(2022\)](#).

1 Data and Summary Statistics

My analysis relies primarily on two data sources. The first is employer–employee matched wage data from the US Census Bureau’s Longitudinal Employer Household Dynamics Database (LEHD). I link the firm identifiers in the LEHD to financial information in the Center for Research in Securities Prices (CRSP)/Compustat merged database obtained from Wharton Research Data Services. I describe the LEHD and CRSP/Compustat–LEHD merged data and samples below. Besides stock return and market cap data, which are from CRSP, all financial variables are from Compustat. Henceforth I refer to the CRSP/Compustat merged sample as the Compustat sample.

1.1 Employee-Employer Matched Wage Data from the LEHD

The Longitudinal Employer Household Dynamics Database (LEHD) contains restricted-use microdata with wage and employer information for individuals in the United States. Wage data in the LEHD are collected from firms’ unemployment insurance filings, and they contain all forms of compensation that are immediately taxable as income, including bonuses and non-qualified stock options. Individuals in the LEHD are linked to their employers through their State Employer Identification Number (SEIN). The LEHD provides crosswalks between the SEIN and the federal Employer Identification Number (EIN), which is also available for firm-level data sources such as Compustat. The LEHD data begin in 1990, although most states join later as the LEHD coverage becomes more comprehensive; the LEHD covers the majority of jobs in the United States by the mid- to late-1990s, and my coverage ends in 2015.¹⁰ The wage data are reported on a quarterly basis, and cover nearly 100 percent of private employees in state-quarters where the data are available. See [Abowd, Stephens, Vilhuber, Andersson, McKinney, Roemer, and Woodcock \(2009\)](#) for a more detailed overview of the construction of the LEHD.

¹⁰An overview of LEHD coverage is provided by [Vilhuber \(2018\)](#).

1.2 Matched Compustat-LEHD Sample

I link firm identifiers in each year of the the LEHD to Compustat records using an internal Census table for mapping LEHD identifiers to Compustat gvkey identifiers. I then assign individuals to their corresponding Compustat gvkey and retain worker-years in which at least one of the worker’s employers was linked to a Compustat firm. Following [Sorkin \(2018\)](#), I convert quarterly LEHD wages to their full-year equivalents. Adjusting wages for a given year following the [Sorkin \(2018\)](#) procedure requires information on wages in the year before and the year after the current year. See appendix section [A](#) for more details. Given the availability of the LEHD wage data from 1990-2015, this means my effective sample period spans 1991-2014. In order to be in the sample, I require a Compustat firm to have at least 15 workers for whom that firm is their primary employer (defined to be the firm where the worker earned the most income that year). I additionally exclude financial firms and regulated utilities from my analysis following common practice. Because I use NAICS industry codes throughout, this excludes the 2-digit NAICS codes 22, 52, and 53.

In appendix Figure [IA.1](#), I show the shares of employment, market cap, and sales represented in my LEHD-Compustat matched sample. On average, I can match LEHD records to Compustat firms covering about 62% of employment, 63% of market cap, and 50% of sales represented in the entirety of Compustat in a given year. These fractions are quite stable over time (panel A of Figure [IA.1](#); since I drop NAICS codes 22, 52, and 53, in panel B of Figure [IA.1](#) I show the shares of Compustat excluding these industries. Among this sample, I cover 71% of employment, 79% of market cap, and 62% of sales among the industries included, which shares are also quite stable over time. I compare the distribution of firm characteristics for my matched sample and the overall Compustat database in appendix Table [IA.1](#). Not surprisingly, firms in my matched sample skew a little bit larger relative to the mean or median firm in Compustat based on assets, employment, or market cap. The distributions of annual excess firm-level stock returns are about the same across the two samples. I obtain the risk-free rate from Ken French’s data library in order to calculate excess returns.

1.3 Variable Construction

My proxy for labor productivity is given by log value-added per worker, following a large literature in labor economics (see [Card, Cardoso, Heining, and Kline, 2018](#), for a review). Firms can have higher labor productivity by enjoying higher total-factor productivity, being more capital intensive, or by hiring more skilled workers. I define value-added for Compustat firms following [Donangelo et al. \(2019\)](#), as the sum of operating income before depreciation, changes in inventories, and labor expenses. Firm-level labor shares are the ratio of labor expenses to value added. In order to estimate the skill level of individual workers (or the average skill of a given firm), I follow a long literature starting with [Abowd et al. \(1999\)](#)—from

now on AKM—in decomposing observed wages into worker- and firm-specific heterogeneity. I start with a modification of the AKM decomposition proposed by [Lachowska, Mas, Saggio, and Woodbury \(2020\)](#) and [Engbom and Moser \(2020\)](#) that allows for the firm-specific component of wages to vary by time. I use LEHD data to create measures of firm-level wage offers and total labor expenses. Details on my construction of all these variables can be found in appendix section [A.1](#).

2 Estimation Strategy

This section develops the identification argument, introduces the empirical specification, and presents evidence that firm-specific stock returns contain powerful and unforecastable variation in labor demand. I then show that the scope for labor supply shocks to contaminate that variation is quantitatively small, and close with the properties of the estimator that bear on interpreting the elasticity estimates.

2.1 Firm-Specific Shocks and Identification of Labor Supply Elasticities in a Static Setting

The firm-specific elasticity of supply is the key quantity that determines wage markdowns in standard models of monopsony with wage posting. To see this, consider a static monopsony framework, which goes back to ([Robinson, 1969](#)). The firm has a concave (net) revenue of labor function $F(L; \Theta)$, where revenues are net of any non-wage labor-related overhead expenses, and Θ is a vector of parameters which determine the mapping of labor L to net revenues. The firm additionally faces, and importantly internalizes, the (inverse) firm-specific labor supply curve $W(L; \Gamma)$, where Γ is a vector of parameters that affects the firm’s labor supply curve, and W is strictly increasing in L . The firm solves:

$$V(\Theta, \Gamma) = \max_L F(L; \Theta) - W(L; \Gamma)L \quad (1)$$

First order conditions give

$$W(L^*; \Gamma) = \frac{\epsilon(L^*; \Theta, \Gamma)}{1 + \epsilon(L^*; \Theta, \Gamma)} F_L(L^*; \Theta) \quad (2)$$

where

$$\epsilon(L^*; \Theta, \Gamma) = \frac{\partial L^*(\Theta, \Gamma)}{\partial W(L^*; \Gamma)} \frac{W(L^*; \Gamma)}{L^*(\Theta, \Gamma)} \quad (3)$$

is the firm-specific labor supply elasticity at the optimal choice of L^* . Note that L^* is directly a function of both Θ and Γ , while W^* is a function of Θ only indirectly through its dependence on L^* . Going forward I suppress the L^* notation and assume any comparative statics are

evaluated at the optimal choices of L and W . Define

$$\frac{\epsilon}{1 + \epsilon} \equiv \text{wage markdown.} \quad (4)$$

The wage markdown represents the fraction of labor's marginal revenue productivity that workers receive in wages, and is determined by the elasticity of firm-specific labor supply. In perfect competition $\epsilon \rightarrow \infty$, and the wage equals F_L , the marginal revenue product of labor; a primary objective of this paper is to quantify the financial value that firms derive from the size of the gap between the wage and the marginal product of labor.

In general, to estimate the labor supply elasticity one can use a shock to any element of Θ , the set of parameters affecting the firm's net revenue of labor function. Meanwhile, shocks to Γ move the labor supply curve and bias supply elasticity estimates. To see this, first consider the estimate of ϵ obtained from a shock to some parameter $\theta \in \Theta$:

$$\hat{\epsilon}_\theta = \frac{\partial L}{\partial \theta} \times \left[\frac{\partial W(L; \Gamma)}{\partial \theta} \right]^{-1} \times \frac{W}{L} = \frac{\partial L}{\partial \theta} \times \left[\frac{\partial W(L; \Gamma)}{\partial L} \times \frac{\partial L}{\partial \theta} \right]^{-1} \times \frac{W}{L} = \frac{\partial L}{\partial W} \times \frac{W}{L} = \epsilon \quad (5)$$

However, an estimate of the elasticity based off some shock to a parameter $\gamma \in \Gamma$ yields a biased estimate:

$$\hat{\epsilon}_\gamma = \frac{\partial L}{\partial \gamma} \times \left[\frac{\partial W(L; \Gamma)}{\partial \gamma} \right]^{-1} \times \frac{W}{L} = \frac{\partial L}{\partial \gamma} \times \left[\frac{\partial W(L; \Gamma)}{\partial L} \times \frac{\partial L}{\partial \gamma} + \frac{\partial W}{\partial \gamma} \right]^{-1} \times \frac{W}{L} \neq \epsilon \quad (6)$$

The bias in (6) stems from the fact that W is both a direct function of γ , and also an indirect function through its dependence on L . Meanwhile, W is a function of θ only indirectly through its dependence on L , so that (5) lacks the additional $\partial W/\partial \gamma$ term that is the source of the bias. This is just a re-statement of the well-known econometric principle that a pure demand shifter identifies the slope of the supply curve. Therefore, in order to estimate ϵ , one needs an instrument whose variation is driven by shocks to Θ , the shifters of the (net) revenue product of labor. Importantly, this means that *any* shock purely to an element of Θ , no matter the source, provides valid identifying variation for estimating the firm-specific supply elasticity. It does not matter whether the shock shifts the slope or the level of the labor demand curve, or if the shock moves labor demand in a negative or positive direction.

2.2 Stock Returns As a Labor Demand Shock

In this section, I lay out the rationale behind my use of firm-specific stock returns to identify supply elasticities. Consider now the firm's value function $V(\Theta, \Gamma)$ which results from optimally solving (1). A shock to some $\theta \in \Theta$ clearly moves firm value, and so when firm value moves because of a shock to θ , we can recover the true elasticity by using the

change in V as the instrument:

$$\hat{\epsilon}_{V,\theta} = \frac{\partial L}{\partial V} \times \left[\frac{\partial W}{\partial V} \right]^{-1} \times \frac{W}{L} = \frac{\partial L}{\partial \theta} \times \left(\frac{\partial V}{\partial \theta} \right)^{-1} \times \left[\frac{\partial W}{\partial \theta} \times \left(\frac{\partial V}{\partial \theta} \right)^{-1} \right]^{-1} \times \frac{W}{L} = \hat{\epsilon}_\theta = \epsilon, \quad (7)$$

since the $\frac{\partial V}{\partial \theta}$ terms cancel out. A similar argument will show that the estimate using firm value shocks driven by some $\gamma \in \Gamma$ leads to the same bias as in (6).

There are some practical advantages to this approach. For example, movements in θ may not be directly unobservable, while firm value is readily available for all public firms. What's more, because firm value is affected by all components of Θ , this makes it a highly-powerful firm labor demand shifter, as its variation simultaneously includes the impact of all the various different labor demand components.

2.2.1 Example labor market model

To make this concrete, consider the following specific formulation of the firm problem (1). The firm faces an upward-sloping labor supply curve, which I assume following [Card et al. \(2018\)](#) takes the form $W(L) = b + \frac{1}{a}L^{1/\beta}$.¹¹ The firm also has market power in the product market, and faces inverse demand curve $P(Q) = \delta Q^{-\phi}$. Firm output is given by $Q(L) = AL^\alpha$. Finally, the firm must pay some overhead cost $c(L) = \kappa L^\zeta$ for every unit of labor it uses. The firm problem is then:

$$\max_L \delta A^{1-\phi} L^{(1-\phi)\alpha} - \kappa L^\zeta - \left(b + \frac{1}{a}L^{1/\beta} \right) L \quad (8)$$

Mapping this to (1), the net revenue function $F(L; \Theta) = \delta A^{1-\phi} L^{(1-\phi)\alpha} - \kappa L^\zeta$, and the parameter vector $\Theta = [A, \delta, \phi, \alpha, \kappa, \zeta]$. Hence, by equation (5), local shocks to *any* of these parameters identify the elasticity of the supply curve at the firm's optimum.

This makes clear that drivers of Θ which serve to identify the supply elasticity could come from a host of sources: pure total factor productivity shifts (driven by A), firm-level product demand shifts (δ), technological shocks that change the production function itself (α), or even changes in firm product market power (ϕ). They could also come from labor overhead cost shifts, so long as they don't move the supply curve directly (κ and ζ). Additionally, while I have not explicitly included capital, one can think of the model as optimizing over labor after optimal choices of other inputs have been made. In this interpretation, capital K and pure total factor productivity are both included in the constant A . Therefore shocks that change capital utilization also provide valid identifying variation.

¹¹This functional form implies a decreasing supply elasticity as firm labor productivity increases, which will turn out to be consistent with the data.

2.2.2 Local projections for stock returns and labor demand responses

With the above discussion in mind, my specification to obtain the baseline average supply elasticity estimate is as follows:

$$\log(Y_{j,t+1}) - \log(Y_{j,t}) = \alpha + \alpha_{I(j),t} + \beta \text{Stock Ret}_{j,t \rightarrow t+1} + \Psi X_{j,t} + \epsilon_{j,t} \quad (9)$$

Here Y = either firm Compustat employment or firm average full-time full-year equivalent adjusted wage.¹² One concern with Compustat employment is that it includes all employees, including those in non-domestic establishments; in section 3.4 I show that using employment from US administrative data in the Longitudinal Business Database yields quantitatively similar patterns in supply elasticity estimates. Elasticities using the LBD are somewhat smaller, meaning my choice to use Compustat employment yields conservative wage mark-downs. The annual stock returns in (9) are given by the sum of the firm’s monthly returns in excess of the risk-free rate from the start of July of year t through the end of June of year $t + 1$. Since firm average wages are accumulated throughout the calendar year, I choose this timing to center the one-year stock returns at the same point within the two adjacent years. The ratio of the estimates $\hat{\beta}^{Emp}/\hat{\beta}^{Wage}$ gives the supply elasticity. Equivalently, this is the two-stage least squares estimate from regressing employment growth on wage growth instrumented by the stock return, which is how I compute standard errors. The $\alpha_{I(j),t}$ denote 3-digit NAICS industry $\times$ year fixed effects, which control for common market shocks. These fixed effects absorb shocks common to a firm’s labor market, which internet appendix IA.3 shows is where an important component of supply-curve-level variation is likely to reside. Section 2.3 states the remaining restriction on firm-specific labor supply shocks precisely and assesses it empirically, after I present the local projection evidence.

The controls $X_{j,t}$, with coefficient vector Ψ , include the contemporaneous change in log firm average AKM worker effects $s_{j,t}$ (defined in equation (A.8) in appendix section A.1); and, lagged growth rates in employment, wages, and firm assets. I include the contemporaneous changes in average worker effects in case stock returns lead to changes in the skill composition of the workforce (for example due to technological change which affects relative skill demands). If this is the case, then failing to account for skill changes may cause employment and wage growth to be misstated in terms of constant efficiency units of labor. Estimates are not sensitive to this control, because stock returns do not appear to induce large skill composition changes on average, as section 2.2.3 shows; section 3.1 relaxes the homogeneous parameter restriction directly by estimating elasticities separately by skill group. My ability to control for such composition changes derived from individual-level LEHD data is an advantage relative to prior firm-level analyses of labor demand responses to firm-specific shocks, such as

¹²The log of the average wage is the relevant quantity for the firm rather than the average log wage, because the average wage is what the firm pays for a single efficiency unit of labor.

Abowd and Lemieux (1993), or more recently Berger et al. (2022). Because of the July t to June $t + 1$ stock return timing, there is some mechanical overlap between the period over which returns are measured and the previous one-year growth rates in wages and employment. I therefore control for the prior one-year growth rates in firm wages and employment. I also control for lagged growth rates in total firm assets, which is known to have some predictive power for prior stock returns; I follow convention for the asset growth predictor by lagging the variable by an additional year (Fama and French, 2015). Due to some extreme outliers in the tails of monthly returns, I cross-sectionally winsorize the outer 0.25% from each tail before summing excess returns to the annual level, although my findings are not meaningfully affected in any way by this decision. The regression requires growth rates in past and future wages and employment and lagged asset growth to all be observable. This means my main analysis sample size is a bit smaller relative to my full matched sample at about 45,500 observations (rounded for Census disclosure purposes).

There are a few advantages of using stock returns. Because they reflect a revision in expected discounted future cash flows, stock returns incorporate the marginal impact of new information about the firm and are largely unpredictable. Additionally, most of the variation in firm-specific stock returns is idiosyncratic. These features make stock returns a potentially powerful and unforecastable labor demand shifter. The preceding argument explains why returns can contain valid identifying variation, but two empirical questions remain. First, do firm-specific returns in fact behave like powerful and unforecastable shocks to labor demand? Second, is the relationship materially contaminated by labor supply shifts that also affect firm value? I address the first question using local projections and then turn to the identifying restriction.

To show that firm-specific stock returns do behave as a powerful and essentially unpredictable labor demand shock, I now study local projections of wages and employment responses over different forward- and backward-looking horizons, which will reveal pre-trends in the relationship between stock returns and labor demand. Similar to Kogan et al. (2020), I examine the growth in average wages or employment over different forward- and backward-looking horizons of h years:

$$\log \left(\frac{1}{|h|} \sum_{k=1}^{|h|} Y_{j,t+k \times \text{sign}(h)} \right) - \log(Y_{j,t}) = \alpha + \alpha_{I(j),t} + \beta_h \text{Stock Ret}_{j,t \rightarrow t+1} + \Psi X_{j,t} + \epsilon_{j,t} \quad (10)$$

Here Y = firm employment or average wage. This specification gives the growth of average employment or wages over the horizon and highlights persistent changes to wages and employment induced by stock return shocks. I estimate (10) for horizons $h = -5$ to $h = 5$ years. For $h = 1$ the specification (10) is the same as (9). Note also that for $h = -1$ the coefficients are set to zero by virtue of the lagged growth controls, and that the $h = 0$ horizon

is zero by definition. I plot the results in the top two panels of Figure 1. Because of the unforeseen nature of a shock to firm-specific stock returns, the pre-trends are non-existent for both wages and employment. Meanwhile, both employment and wages display large, statistically significant, and persistently positive responses at the time of and following the stock return shock. The joint positive employment and wage responses to movements in stock prices support that stock returns constitute a powerful and essentially unpredictable labor demand shock.

In the bottom two panels of Figure 1 I examine the relationship with firm productivity outcomes. In particular, I estimate (10) for the growth in average firm total factor productivity (from İmrohoroglu and Tüzel (2014)) or the growth in average labor productivity (value-added per worker). Consistent with stock returns containing unpredictable and powerful immediate information about firm productivity, I find a highly significant positive response of both TFP growth and labor productivity growth, with little to no pre-trends.¹³ This affirms that stock returns do contain information about labor demand determinants like firm productivity. To underscore the point, I further estimate (10) with total capital stock growth as the dependent variable in appendix Figure IA.2. Again I find a starkly positive response starting at the initial one-year horizon, with no pre-trends.

The dynamic patterns found in Figure 1 are also of note. First, observe that wages and productivity outcomes all immediately increase and then subsequently mean-revert, while employment continually increases over the 5-year horizon. This slow adjustment of employment is consistent with the presence of adjustment costs. A similar dynamic pattern, with even slower adjustment, can be found for the capital response in appendix Figure IA.2. Meanwhile, the wage response is initially high and subsequently declines. In section 4, I show that a calibrated dynamic framework with labor adjustment costs and gradual accumulation of labor can match these patterns. Importantly, the labor adjustment costs mitigate the firm’s market power relative to the static first-order condition (2).

2.2.3 Heterogeneity in local projections by worker skill

Having demonstrated evidence of stock returns’ relevance for labor demand, with a lack of pre-trends in local projections, I now relax the assumption that the firm hires a single labor type. The homogeneous setup imposes the same labor market parameters for all skill groups, but in practice workers of varying skill levels may face different labor market frictions, and consequently exhibit more or less elastic labor supply. I therefore extend the local projection procedure to analyze labor demand responses separately by worker skill type. To proxy for skill level, I group individual workers into low-, middle-, and high-skill workers—defined as being in the bottom two quintiles; third and fourth quintiles; or top quintile in the cross-

¹³One of the pre-period coefficients is significant for labor productivity, though the economic magnitude is small. Between the specifications in Figure 1, appendix Figure IA.2, and Figure 2, there are a total of 44 pre-period coefficients, so it isn’t surprising to find a significant coefficient by chance.

sectional distribution of worker-level intercepts in the AKM wage decomposition, respectively. I re-estimate equation (10) using firm-level average wages and employment for each skill group. I also modify the firm skill composition control to be the within skill-group average. Since I can't directly observe the Compustat employment by worker skill type, I assume the number of workers of a given skill level are proportional to their observed shares in the LEHD. So if the firm has $N_{j,t}$ workers in the sample of the AKM decomposition (A.5) in year t , $N_{j,t}^k$ of which are of skill group k , then I assume Compustat employment in skill group k is equal to $EMP_{j,t}^k = EMP_{j,t} \times N_{j,t}^k / N_{j,t}$.¹⁴

Estimates are displayed in Figure 2. Three key patterns emerge from the figure. First, in the top panel, the skill-specific employment responses are close to proportional for each skill group, and are quite similar to the employment responses in homogeneous specifications found in Figure 1. This suggests that stock returns do not entail large composition changes within the firm on average. Second, while there are highly-significant positive and immediate wage increases which subsequently decline across all worker skill levels, the magnitude of the wage response is much larger for the top skill group. This points to highly-skilled workers exhibiting less elastic labor supply, implying larger labor market frictions for this group. However, as I will show in section 5, dynamic adjustments which mitigate static labor market power are especially prominent among the most skilled workers. This means skilled worker wage markdowns are somewhat smaller than supply elasticities would suggest. I return to analyze the implications of this skill heterogeneity in detail in sections 3 and 5.

2.3 Quantitative Importance of Labor Supply Shocks

Taken together, the local projection results show that firm-specific stock returns contain substantial, unforecastable variation in firms' desired labor demand. This evidence does not by itself establish the exclusion restriction. It does, however, narrow the remaining question: whether stock returns also contain sufficiently important labor supply variation to materially affect the estimated ratio of employment and wage responses.

Returning to the example labor market model (8), the vector of supply-shifting parameters is given by $\Gamma = [b, a, \beta]$. Hence shocks to b , a , or β shift the supply curve and do not provide valid identifying variation. As Card et al. (2018) explain, these parameters have clear economic interpretations. The parameter b maps into workers' outside option, β captures the ease of substitution across firms and is closely related to workers' bargaining power, and a controls the level of the supply curve, including through workers' perceptions of firm-specific amenities.¹⁵

¹⁴I continue to use Compustat employment because Compustat reflects workers at all the firm's establishments instead of just those that are covered by the LEHD in that year. This is especially helpful in the first part of my sample when the LEHD doesn't cover all states. As I show in section 3.4, Compustat employment is also quite a bit more responsive to stock return shocks, leading to higher—and hence more conservative—supply elasticity estimates.

¹⁵Two of the three are largely outside the scope for firm-specific, high-frequency variation. Essentially all structural labor market models treat β as a deep parameter, not subject to idiosyncratic firm-specific shocks

The relevant identifying assumption is not that firm-specific labor supply shocks never occur. Rather, conditional on market-level controls and other covariates, such shocks must not account for enough of the return-induced wage and employment variation to materially distort the estimated supply elasticity.

Several pieces of evidence indicate that supply-side contamination is quantitatively limited. First, internet appendix [IA.3](#) shows that an important component of the level of the labor supply curve is driven by market-wide wages and productivity. Given that firm-specific stock returns contain both systematic and idiosyncratic components, this is a potentially important source of confounding variation. Depending on the model of labor markets, this variation can be absorbed with different combinations of market-by-year fixed effects and other controls for shocks to labor market competitors. Appendix Table [IA.3](#) reports nine such specifications, spanning alternative labor market definitions, competitor wage and employment growth, competitor stock returns, and local unemployment and concentration, and both the average elasticity and the gap between high- and low-productivity firms are unmoved across all of them. Adding every one of these controls at once gives elasticities of 3.564 (s.e. = 0.576) and 1.727 (s.e. = 0.182) for low- versus high-productivity firms, against baseline estimates of 3.525 (s.e. = 0.574) and 1.683 (s.e. = 0.173). Appendix Table [IA.4](#) reaches the same conclusion from a different direction, using within-firm, cross-establishment variation and industry $\times$ commuting zone $\times$ year fixed effects; because those specifications rely on LEHD rather than Compustat employment the levels are lower, but the productivity spread survives.

Second, pervasive cross-sectional differences in fixed firm-specific amenities are accounted for by estimating supply elasticities in log differences rather than levels; only high-frequency within-firm changes in amenities can bias the estimates ([Sorkin, 2018](#); [Lamadon et al., 2022](#); [Sorkin, 2024](#)).¹⁶ I also control directly for observable firm-specific labor supply shocks. Appendix Table [IA.11](#) shows that controlling for unionization events and changes in non-compete enforceability, entered both as levels and interacted with the return, leaves the elasticity estimates essentially unchanged, even in samples concentrated among firms that experience these events. Variation in hours provides another observable change in job characteristics and labor input that could affect measured wages and employment. In

and particularly not at annual frequency; see [Card et al. \(2018\)](#), [Lamadon et al. \(2022\)](#), and [Kroft et al. \(2025\)](#), or [Stole and Zwiebel \(1996\)](#), [Acemoglu and Hawkins \(2014\)](#), and [Kuehn et al. \(2017\)](#) for bargaining models with closely-related structural parameters. The reservation wage b , being an outside option, is likewise naturally determined at the market rather than the firm level. Both therefore fall largely within what the market $\times$ year controls absorb, which leaves a as the parameter of practical concern.

¹⁶Two further reasons favor an instrumental variables estimate over ordinary least squares even in the presence of such shocks. First, as internet appendix [IA.3](#) shows, the bias in the IV estimate can be signed, whereas OLS elasticity estimates can be biased in either direction for any volatility of firm-specific amenities shocks. Second, there is likely some measurement error in firms' wages—from imperfections in the linkage between LEHD wage information and Compustat financial information, and from incomplete LEHD coverage of firms' establishments—which would attenuate OLS estimates toward zero and yield overly large wage markdowns, while an instrumental variables estimate corrects for it.

appendix Table [IA.2](#), hours-adjusted estimates—which express employment in total hours and wages per hour, using usual hours for LEHD individuals matched to the American Community Survey—are close to the baseline estimates across the productivity distribution for the same sample, as are estimates that hold the set of observable LEHD states fixed and estimates that control for lagged and contemporaneous markup growth.

Third, section [3.2](#) re-estimates the elasticity using three alternative labor demand instruments in place of own-firm stock returns: customer stock returns, patent-induced shocks, and stock returns in narrow windows around earnings announcements. Each rests on different identifying variation, and each reproduces the elasticity gap between high- and low-productivity firms, though less precisely than the baseline. Overidentified specifications that combine each alternative with the baseline instrument yield similar estimates and pass their Hansen J -tests. For a supply-side confound to be generating these results, it would have to contaminate four instruments drawing on different information in the same direction and by close to the same magnitude. No single exercise rules out every possible firm-specific supply shock. Taken together, however, the stability of the estimates across market controls, observable supply shocks, hours adjustments, and distinct instruments indicates that such variation is unlikely to be quantitatively important for my findings.

Finally, internet appendix [IA.3](#) characterizes the direction of any remaining bias. Failing to account for market-wide shocks is likely to bias elasticity estimates downward, causing me to make erroneously large inferences about labor market power. Once common shocks are accounted for, if firms do not cut amenities by too much on average when idiosyncratic productivity improves, residual amenities shocks bias the estimated elasticity upward. Thus any remaining bias would make the implied labor market power estimates conservative. Exact identification requires the stronger condition that only market-specific shocks move the level of the supply curve, which is imposed in essentially all work estimating the passthrough of firm-specific shocks to workers.¹⁷ I view this result as a characterization of the residual bias that the empirical exercises cannot eliminate, rather than as the primary support for the estimation strategy.

2.4 Properties of the estimator

Three properties of the estimator remain to be established before turning to the elasticity estimates: that noise in firm value does not bias it, that the argument extends to dynamic labor choice, and how it behaves under shocks that move firm value and labor demand in

¹⁷See [Lamadon et al. \(2022\)](#), [Kline et al. \(2019\)](#), and [Garin and Silverio \(2020\)](#), for example. The assumption is necessary in that work because unobservable firm-specific amenities shocks bias passthrough coefficients. Papers in the asset pricing literature that incorporate labor market frictions, such as [Kuehn et al. \(2017\)](#), [Belo et al. \(2014\)](#), and [Donangelo \(2014\)](#), likewise assume labor supply shocks are determined at the market level. The condition must be invoked even when the identifying variation is quasi-experimental—as in [Lamadon et al. \(2022\)](#) and [Kroft et al. \(2025\)](#), which both use variation from winning construction lotteries—because otherwise firms could adjust their amenities in response to the exogenous shock.

opposite directions. I take these in turn, and close by relating the resulting markdowns to a bargaining protocol.

2.4.1 Noise in firm value/stock returns

While firm value (and hence the stock return) contains information that is informative about labor demand, it also could contain information that is unrelated to labor demand or supply. This does not bias estimates, although it does naturally make the estimate noisier. To illustrate the point, consider the following modification on (1):

$$V(\Theta, \Gamma, \Lambda) = G(\Lambda) + \max_L [F(L; \Theta) - W(L; \Gamma)L] \quad (11)$$

The function $G(\Lambda)$ is some production technology which generates value to the firm but is unrelated to today's labor demand decision. Now, suppose that firm value is affected by a shock to some $\lambda \in \Lambda$ and some $\theta \in \Theta$, and we estimate the supply elasticity via the total change in firm value. A simple derivation in appendix IA.4 shows that the elasticity estimate recovered from the resulting total shock to firm value is identical to the one in (5), and hence yields the correct estimate. The intuition is that the noise reduces the marginal effects of firm value on employment and wages proportionally, so that the elasticity itself (which is the ratio of the two effects) is unaffected.

When labor supply decisions are made in a static matter, one obvious candidate for such “noise” in stock returns is any information about future firm outcomes, since stock returns incorporate both information about the near term and the future. For example, consider the value of an infinitely-lived firm with static labor demand:

$$V_0 = E_t \left[\sum_{t=0}^{\infty} D_t(\Theta_t, \Gamma_t) / (1 + r)^t \right] \quad (12)$$

where $D_t(\Theta_t, \Gamma_t) = \max_{L_t} [F(L_t; \Theta_t) - W(L_t; \Gamma_t)L_t]$. This demonstrates that if the choice of L_t is made in a static fashion—a common assumption in many models—one can think of all expectations about any future quantities as being embedded in the $G(\Lambda)$ term. This includes information about firm discount rates, future productivity shocks, and so on. Accordingly, under a static view of firm-level labor determination, the impact of discount rate shocks and any other future expectations serve to reduce precision in my estimates, but do not cause bias. Note that this is also only the case if the information about the future contains no information about the present. For example, in a model with persistent idiosyncratic TFP shocks, an improvement in firm-level TFP today clearly increases labor demand today, even though it also improves future productivity expectations.

2.4.2 Extending to dynamic settings

In reality, it is likely that labor demand is not a fully static decision for firms. In a model with labor adjustment costs, news about future profits—for example future productivity or demand shocks—can cause the firm to optimally adjust employment now. In fact, this is true even for pure discount rate shocks, which cause the firm to optimally re-allocate output intertemporally, affecting labor demand immediately. Such shocks can be incorporated into the basic static framework in (1) by appropriately expanding the parameter space: any relevant future information for today’s labor demand decision simply becomes embedded in the Θ vector, meaning these shocks also provide valid identifying variation.

While information about the future can provide valid identifying variation for supply elasticities in a dynamic setting, the mapping from supply elasticities to wage markdowns may be somewhat different than (4). I expound on this idea further in section 4, where I explicitly write out a dynamic monopsony framework and show how to use it to map supply elasticity estimates and other empirically-observable quantities into wage markdowns. The incentive to avoid future labor adjustment costs can cause the firm to curtail wage markdowns, relative to what would be implied by empirical supply elasticity estimates in a purely static setting. This turns out to be quantitatively important. Without making such adjustments, the implied wage markdowns are still directionally correct, but the implied magnitudes of wage markdowns are quite large. This makes it harder to reconcile with time series patterns in labor shares and cross-sectional differences in firm profitability.

2.4.3 Demand shocks that comove negatively with stock returns (e.g. markups and production technology shifts)

Some parameters in Θ may increase firm value while *decreasing* labor demand. As made clear above, whether a shock to revenue productivity moves demand up or down (regardless of the direction it moves firm value) does not matter for identifying the supply curve. However, a practical challenge for estimation is that such shocks can reduce the power to detect an effect. To take an extreme example, if half of the variance in stock returns were driven by shocks to parameters that move labor demand down but increase firm value, and the other half were driven by demand shifters that move both firm value and labor demand in the same direction, then a regression of employment or wage growth on stock returns would exhibit zero relationship. In other words, this is not an issue of endogeneity, but of instrument strength and monotonicity.

Pure product market power shocks are likely a good example of such shocks. In classic models of monopoly, increased market power causes firm profits to increase, but also causes them to reduce their output (Syverson, 2019; De Loecker et al., 2020). This can cause a component of stock returns to correlate negatively with labor demand. Technological shocks

that reduce the output elasticity of labor can work in a similar manner.¹⁸ At the same time, markups covary closely and positively with firm (revenue) productivity (Autor et al., 2020; Pellegrino, 2025). Thus shocks to product market power are likely to co-occur with positive revenue productivity shifts, which tend to push labor demand and firm value in the same direction. Consistent with productivity effects dominating in practice, the local projections above show that firm labor demand and productivity respond positively, immediately, and strongly to firm-specific stock returns.

However, an important subtlety arises when considering for *which* firms I can identify the slope of the labor supply curve. Because labor demand tends to move positively with stock returns on average, I will identify the supply curve in firms (or during time periods) where stock returns are dominated by positively correlated labor demand shifters like productivity shocks. In contrast, I will have less power to identify the slope of the supply curve in firms (or time periods) where returns are driven primarily by shocks such as changes in markups. This could matter if the labor supply curve has heterogeneous slopes in the cross section of firms or in the time series—and I will show both are indeed the case. It is worth being precise about the nature of this concern. Opposite-signed demand shocks do not violate the exclusion restriction; what they can violate is monotonicity, in the sense that the first stage may be weak or wrong-signed for particular firms or periods. The consequence is therefore not bias but reweighting: the estimator recovers the supply elasticity where returns are dominated by same-signed demand variation, weighting firms by their first-stage strength. That has a testable implication. If first-stage strength varied across firms in a way that materially changed whose supply curve I identify, then the homogeneous estimate—which weights firms by first-stage strength—should diverge from an equal-weighted average of the disaggregated estimates I report in section 3. It does not. Averaging the productivity-quartile estimates reported in section 3.2 gives 2.64, the above- versus below-median split reported in the same section gives 2.60, and averaging the two subperiod estimates in section 3.3 gives 2.40, against a homogeneous estimate of 2.53. Heterogeneity in first-stage strength across firms therefore does not appear to be quantitatively important for the elasticities I report. Consistent with this, I also explore specifications which explicitly control for markup shocks, and adding such controls has no impact on my empirical estimates.

2.4.4 Implications under bargaining protocols instead of wage-posting

Besides monopsonistic wage-posting models, other models of imperfect competition in labor markets imply that firm-specific shocks to marginal productivity should affect wages and employment. In internet appendix section IA.5, I explore the interpretation of firm-specific

¹⁸This can be easily seen by examining the exponent in equation (8): the elasticity of firm product demand parameter (ϕ) and the Cobb-Douglas output exponent (α) get multiplied together in the firm’s revenue function. Thus the Cobb-Douglas exponent of labor in the revenue production function is a combination of these two forces, and shocks to either behave similarly.

shocks on labor demand when wages are instead determined by multilateral bargaining in multi-worker firms, as in [Stole and Zwiebel \(1996\)](#) and [Acemoglu and Hawkins \(2014\)](#), and subject to convex hiring costs, as in [Garin and Silverio \(2020\)](#). I show that the wage markdowns from this framework are closely related to the wage markdowns that are implied by wage posting. See internet appendix section [IA.5](#) for further details.

3 Main Supply Elasticity Estimates

Having established the identifying framework and empirical specification in section 2, I now use equation (9) to estimate labor supply elasticities for the pooled sample and examine heterogeneity by worker skill, firm labor productivity, and time.

3.1 Homogeneous and by worker skill

I report estimates of equation (9) for all workers and by worker skill type in Table 1, as well as the implied supply elasticities. The one-year wage and employment coefficients in the table correspond to the $h = 1$ response coefficients plotted in Figures 1 and 2; their ratio yields the estimated labor supply elasticity. Looking at the first column, I find an estimate of 0.12 for employment growth in the first row, and 0.046 for wage growth in the second row. In the third row I take the ratio of the two estimates to arrive at my baseline pooled average supply elasticity estimate of about 2.53 (standard error of 0.3). This is in line with the range of estimates in previous research, although on the slightly smaller end.¹⁹ This is likely because I focus on large publicly-traded firms.

The next three columns of Table 1 explore estimates of (9), modified to focus on heterogeneity by worker skill, as in section 2.2.3. I find larger wage responses for the most skilled workers, which translates into a smaller supply elasticity: the elasticity is 1.2 (standard error of 0.17) for the most skilled workers versus 5.9 (standard error of 0.59) for the middle skill group, and 7.97 (standard error of 0.88) for the lowest skill group. This suggests larger labor market frictions for skilled workers. In section 5, I examine the quantitative implications of this fact for the role of skilled workers in driving firm profits from labor market power.

3.2 Cross-firm heterogeneity in supply elasticities by labor productivity

I now examine heterogeneity in firm-level labor supply elasticities by firm labor productivity, defined by log value-added per worker. Labor productivity is proportional to Cobb-Douglas marginal revenue product, and marginal product dispersion is ruled out in frictionless, perfect competition models. Differences in market power can be a natural

¹⁹[Sokolova and Sorensen \(2021\)](#) find the median supply elasticity estimate is around 1.7 in a meta-analysis. [Bassier et al. \(2022\)](#) point out that recent quasi-experimental evidence—such as [Kroft et al. \(2025\)](#), [Cho \(2018\)](#), and [Dube, Manning, and Naidu \(2018\)](#), and [Caldwell and Oehlsen \(2018\)](#)—has found supply elasticities between 2 and 5.

source for marginal product dispersion (Hsieh and Klenow, 2009; Peters, 2020), making labor productivity a natural dimension on which to test for differences in labor market power. Moreover, stylized facts for productive firms strongly suggest they exhibit more market power: in panel A of Table 2 I examine the relationship between labor productivity, labor shares, valuations, and operating profitability. The table shows that productive firms have lower labor shares, higher operating performance (current or future return on assets), and higher valuation ratios (as measured by Peters and Taylor (2017) intangible-adjusted Tobin’s Q or the market-to-book ratio).

Crouzet and Eberly (2021) and Abel and Eberly (2011) show that the Tobin’s Q valuation ratio includes both market power rents and marginal Q, the marginal increase in the value of the firm with respect to an additional incremental unit of capital. If productive firms had high valuation ratios because of higher marginal Q, one would expect to see them have higher investment rates. Panel B of Table 2 confirms this is not the case: despite strongly increasing valuations across labor productivity quartiles, firms in the top quartile of labor productivity do not have higher significantly higher average investment rates than firms in the bottom quartile. These stylized facts strongly point to empirical labor productivity dispersion correlating with dispersion in economic rents earned from (product or labor) market power.

With these characteristics of productive firms in mind, I estimate a modified version of (9) in Table 3 to test for differences in labor market power between high- and low-labor productivity firms. I first sort firms into quartiles on log value-added per worker, and then I estimate heterogeneous coefficients for each quartile of the productivity distribution:

$$\log(Y_{j,t+1}) - \log(Y_{j,t}) = \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} + \sum_{q=1}^4 \mathbf{1}(q(j,t) = q) \times \beta_q \text{Stock Ret}_{j,t \rightarrow t+1} + \Psi X_{j,t} + \epsilon_{j,t} \quad (13)$$

where again Y = firm Compustat employment or firm average full-time full-year equivalent adjusted wage. Here $q(j,t)$ denotes the productivity quartile of firm j at time t . The controls $X_{j,t}$ are the same as introduced in (9), including lagged employment, wage, and asset growth, and the contemporaneous change in firm average worker skill. Finally, I add quartile-specific fixed effects $\alpha_{q(j,t)}$ to allow each productivity quartile to have a different intercept and slope.

In Table 3 I estimate (13) for the overall firm and for each skill type. Table 3 shows that more productive firms face lower supply elasticities at the firm level and within each skill level. Panel A shows that the overall supply elasticity for a highly-productive firm is about 1.17 (standard error of 0.15), while it is 4.32 (standard error of 0.87) for a low productivity firm. These estimates mask heterogeneity by worker skill level; panel D shows that the highest skill workers have much lower supply elasticities, implying greater wage markdowns are largest for skilled workers at productive firms. However, even for low- and middle-skill workers (in panels B and C, respectively) the elasticities are decreasing in firm productivity. The fifth

column shows that the differences in the supply elasticity estimates between the top and bottom quartile firms are also always statistically significant. These differences are driven primarily by higher wage responses in the denominator—highly-productive firms move wages by more to induce a given change in employment relative to unproductive firms.

These supply elasticity estimates immediately imply a wide gap in rents earned from wage markdowns across the productivity distribution. This in turn should affect the differences in labor shares, valuations, and income to capital owners in the cross-section of firms sorted on labor productivity. In Table 5 I compute the implied static log wage markdown for each firm j at time t based off the productivity quartile-specific supply elasticity

$$\log(\text{Markdown})_{j,t} = \log \left(\frac{\widehat{\epsilon_{q(j,t)}}}{1 + \widehat{\epsilon_{q(j,t)}}} \right), \quad (14)$$

where $q(j, t)$ is firm j 's current labor productivity quartile. Similar to Table 2, I then analyze the relationship between $\log(\text{Markdown})_{j,t}$ and the firm log labor share; total Tobin's Q from Peters and Taylor (2017); and current and future firm profitability. Relative to Table 2, I explore more stringent specifications by adding controls for capital intensity (the log of physical capital and intangible capital per worker), and I also control for imputed firm-level log markups using the method from De Loecker et al. (2020). Additionally, I explore specifications with and without firm fixed effects.

For each specification, I find a highly statistically significant relationship between estimated log wage markdowns and firm outcomes, and always with the expected sign. With the definition of the markdown in (4), a higher log markdown means lower firm rents from market power. Accordingly, the estimated log markdown is associated with higher firm log labor share. Examining the first column, the elasticity of labor share to wage markdown is 0.90 (standard error of 0.11), close to the theoretical value of 1, and highly statistically significant; using within-firm variation in the second column, the elasticity declines to 0.69, but again remains highly significant. In the next 2 columns, I find a robust negative relationship between log markdowns and the Peters and Taylor (2017) total Q valuation ratio. This is consistent with rents from labor market power leading to higher firm valuations. In the last four columns I examine the relationship between the current or future return on the firm's capital assets and log markdowns. I again find highly significant relationships, with the expected negative sign. Comparing the second-row coefficients on log markups further underscores the strong relationship between estimated log wage markdowns and firm outcomes: although log markups consistently exhibit the expected sign and the coefficients are always statistically significant, they are generally smaller and estimated less precisely.

I now report in detail the alternative-instrument estimates summarized in section 2.3. In Table 4, I provide elasticity estimates using alternative labor demand shifters: instead of own-firm stock returns, I use the stock returns from firms' customers as an indirect demand

shock; I leverage the market value of firm patenting following [Kogan, Papanikolaou, Seru, and Stoffman \(2017\)](#); and finally, I use the average stock returns in excess of the market only in the 4-day windows around firms’ quarterly earnings announcements to isolate information about earnings. I discuss the construction and identification assumptions underlying each of these alternative shocks in appendix section [IA.6.3](#).

In Table 4 I split firms by above- and below-median productivity, both for brevity and because the alternate instruments have less statistical power. The top panel uses customer returns to estimate the elasticity. Because customer returns are available for only about a quarter of firm-years, I compare elasticity estimates to my baseline specification within the same subsample. I find elasticity estimates of 3.40 (s.e. = 1.47) for below median productivity, and 1.45 (s.e. = 0.52) for above median productivity; the p-value on the difference is 0.106. My baseline estimates within the same subsample are 2.86 (s.e.=0.55) and 1.47 (s.e.=0.24). In the next panel of Table 4 I use [Kogan et al. \(2017\)](#) patenting and earnings returns, which are both available for the main sample. Both these instruments return higher elasticities for unproductive firms. For the patenting shock, the elasticity estimates are 2.99 (s.e.=0.65) and 1.97 (s.e.=0.69) for unproductive versus productive firms, respectively, and the difference has a p-value of 0.038. For earnings returns, the corresponding estimates are 3.34 (s.e. = 0.66) and 1.87 (s.e = 0.29), with a p-value of 0.011. Meanwhile, baseline estimates within the full sample are 3.53 (s.e. = 0.57) and 1.67 (s.e. = 0.17), with highly significant p-value.

Finally, in the last panel I explore overidentified specifications where I use both the given alternate instrument and the original stock return instrument. In all cases the key cross-sectional spread is highly significant, and close to the baseline. The final column additionally reports the p-value from a Hansen J-test, where I always fail to reject the overidentifying restrictions, which is helpful for instrument validity. Overall, the fact that these alternate instruments also uncover the same economic patterns—albeit with less statistical precision—as my main specification lends credence to my supply elasticity estimation strategy.

3.3 Time-Variation in Labor Supply Elasticities

Besides the cross-sectional spreads in firm valuations and labor shares, time series changes in the labor supply elasticity could speak to changes in the aggregate labor share. To investigate this possibility, I re-estimate supply elasticities via equation (9) both for the homogeneous worker specification and by worker skill for overlapping moving 3-year windows. Figure 3 shows the results. Supply elasticities for all worker skill levels have declined over time, suggesting a secular rise in labor market frictions; however, the aggregate trend (in yellow) is driven primarily by the pattern for the most skilled workers (in green). This suggests the labor market frictions faced by skilled workers may play a particularly important part in the quantitative role labor market power plays in explaining the time series decline in the labor share.

Overall, when I estimate supply elasticities via (9) for the first (1991-2002) and second (2003-2014) halves of my sample period, I find an average labor supply elasticity of about 2.92 (standard error of 0.30) in the 1991-2002 period, and 1.88 (standard error of 0.17) in the 2003-2014 period. In section 5, I explore the quantitative role the time-series labor supply elasticity decline has had in explaining the aggregate labor share decline.

A number of papers have noted that the time series decline in the labor share can be attributed in part to a widening cross-sectional labor share gap (Kehrig and Vincent, 2021; Gouin-Bonenfant, 2022; Hartman-Glaser et al., 2019). I now re-estimate (13) separately for the first half (1991-2002) and second half (2003-2014) of my sample. I plot the elasticity estimates in a bar chart in Figure 4. I find a secular decrease in supply elasticities across the productivity distribution, while the cross-sectional labor productivity ranking in elasticities remains consistent over time.

Comparing the estimates for 1991-2002 to 2003-2014, the supply elasticity for firms in the bottom quartile of the labor productivity distribution changes from 4.86 (standard error of 1.3) to 3.25 (standard error of 0.57), and from 1.34 (standard error of 0.14) to 0.85 (standard error of 0.12). Applying the simple static markdown formula (4) suggests that the cross-sectional gap in implied markdowns widened between the first and second half of my sample. In the 1991–2002 period, productive firms’ markdowns implied that workers received about 26% less of their marginal product than at unproductive firms (1.34/2.34 minus 4.86/5.86). In the 2003–2014 period, the implied gap increased to about 31% (0.85/1.85 minus 3.25/4.25).

With this result in mind, I explore further whether changes in wage markdowns track the cross-sectional gap in log labor shares at a higher frequency. I estimate (13) for overlapping moving 3-year windows and then take the elasticity estimates for the top- and bottom-productivity quartiles. I use the standard markdown formula to estimate the elasticity-implied spread in log markdowns between the two productivity types and compare this to the spread in log labor shares. The log markdown spread is

$$\text{Log Markdown Spread}_t = \log \left(\frac{\widehat{\epsilon}_{4,t}}{1 + \widehat{\epsilon}_{4,t}} \right) - \log \left(\frac{\widehat{\epsilon}_{1,t}}{1 + \widehat{\epsilon}_{1,t}} \right) \quad (15)$$

where $\widehat{\epsilon}_{q,t}$ is the estimated supply elasticity for productivity quartile q in the 3-year moving window centered at time t . At the same time, I compute the average log labor share of the high- and low-productivity firms over the same 3-year window:

$$\text{Log Lshare Spread}_t = (1/3) \sum_{\tau=t-1}^{t+1} (E[\log(lshare_{4,\tau})] - E[\log(lshare_{1,\tau})]) \quad (16)$$

The log labor share is linear in the log wage markdown, all else held constant, so if these estimates reflect genuine differences in market power then they should move together over time.

Figure 5 plots the two standardized series. Due to the availability of LEHD wage data and the requirements of my regression specification, I compute supply elasticity estimates for the years 1992-2013. There is a clear tight link between the cross-sectional dispersion in labor shares and the wage markdowns for firms sorted on labor productivity, with the two series tracking each other closely over time. The correlation is about 0.73, with a Newey-West t-stat of 4.10. Since there are only 22 observations I compute the small-sample adjusted t-stat, and I choose a lag length of 5 years due to the persistence of the two series. The widening gap in labor shares and elasticity-implied markdowns in Figure 5 mirrors the changing relative valuations and productivity levels of high- and low-productivity firms, which I show in appendix Figure IA.5, suggesting the phenomena may be jointly linked.

3.4 Supply Elasticity Robustness Checks

I next provide details on the robustness checks summarized in section 2.3. I cover the empirical specifications and data underlying these exercises in section IA.6 of the internet appendix.

Since a potentially important source of confounding variation comes from aggregate shocks, I carefully check to see how my estimates vary with different controls for common market level shocks in appendix section IA.6.1. The resulting estimates are contained in appendix Table IA.3. I explore both homogeneous specifications and sorts by labor above- and below-median labor productivity rank. Supply elasticity estimates are quantitatively stable across the alternative controls for common market shocks. For example, consider the productivity sorts in panel B. In the last column of panel B, I add a “kitchen sink” set of labor market controls on top of my baseline specification: labor market competitor employment and wage growth; controls for the average growth in local unemployment rates and local labor market concentration; labor market competitor stock returns; and fixed effects for empirically-estimated labor market $\times$ year fixed effects.²⁰ The corresponding supply elasticity estimates are 3.564 (s.e. = 0.576) and 1.727 (s.e. = 0.182) for low- versus high-productivity firms, respectively; these are economically and statistically indistinguishable from the baseline estimates of 3.525 (s.e. = 0.574) and 1.683 (s.e. = 0.173) in the first column. Comparisons of estimates across the remaining columns of the table with different control combinations yield the same basic conclusion. This exercise demonstrates that the method I use for accounting for common market shocks is quantitatively immaterial to my findings; consequently I stick to the baseline specification for parsimony.

In appendix Table IA.4 I explore alternative specifications for accounting for market-level shocks by explicitly exploiting within-firm, cross-establishment variation in exposure to local labor markets (defined as industry $\times$ commuting zone, as in Berger et al. (2022)). In the top panel I sort firms on their labor productivity advantage relative to the average of productivity

²⁰See appendix section IA.6.1 for an explanation of how these variables are constructed.

of firms in the same local market.²¹ In the bottom panel, I alternatively sort firms based on their local market share following [Berger et al. \(2022\)](#). These specifications now include very fine industry $\times$ local commuting zone $\times$ year fixed effects to absorb market-level shocks. Because these specifications rely on LEHD employment, which is less responsive to firm shocks, the elasticity estimates are on average lower, but broadly consistent with my main estimates. Reassuringly, I continue to find a large spread in labor supply elasticities across the productivity distribution when using this method. In the bottom panel I also capture a similar pattern to [Berger et al. \(2022\)](#), in that firms with a larger local wage bill share face lower elasticities. See appendix section [IA.6.2](#) for details.

Next, in appendix Table [IA.2](#), I examine specific concerns related to my measurement. As discussed in section [2.4](#), shocks that positively affect returns but reduce labor demand do contain valid identifying variation. However, they can change which firms or time periods drive the variation that identifies the supply curve, which can affect the point estimates I obtain in practice because supply elasticities are heterogeneous by firm and time. As I explained in section [2.4](#), shocks to markups are a particularly likely candidate for this type of shock. Accordingly, I explore adding controls for lagged and contemporaneous growth in markups in the first row of the table. The first two columns compare the point estimate in the 1991-2002 and 2003-2014 halves of the sample, and the last two columns for high- versus low-productivity firms. When comparing with my baseline estimates in the second row, I find that accounting for changes in markups has no effect on my estimates.

The next measurement concern is that the LEHD doesn't reach its full coverage of states until the end of 2003, and I estimate average wages for a firm each year only across its operations in currently-observable states. In the third row of appendix Table [IA.2](#), I re-estimate specifications using only information from just the 15 states that were in the LEHD by the end of 1991.²² The fourth row uses my baseline procedure with all LEHD states as they become available, except restricted to the overlapping sample (some firms operate only outside these 15 states and consequently drop out when I restrict to these states). Again the difference in supply elasticity estimates in the time series or cross-section is quantitatively minimal relative to my baseline measurement of firm average wages. Thus the changing set of observable LEHD states plays no role in generating my key empirical findings.

Finally, another measurement concern could be that I do not observe hours in the LEHD. Firms could increase employment on the intensive hours margin in the short run, biasing elasticity estimates based off the number of employees. I address this by merging individuals in the LEHD to the American Community Survey (ACS), and I extract each matched person's reported usual working hours. I compute the average hours by firm, requiring firm-years to

²¹Sorting on this dimension locally is also consistent with the model of [Lindenlaub, Oh, and Peters \(2024\)](#), who introduce a labor market search model with firm sorting to locations based on productivity.

²²These states are Maryland, Alaska, Colorado, Idaho, Illinois, Indiana, Kansas, Louisiana, Missouri, Washington, Wisconsin, North Carolina, Oregon, Pennsylvania, and California.

match to at least 20 individuals with observable ACS hours. Since the ACS is only available from 2005 and on, I only report the impact on cross-sectional elasticity patterns in the final two rows of appendix Table [IA.2](#). In the fifth row I estimate hours-adjusted supply elasticities by expressing employment growth in total hours, and wage growth in per hour units. Relative to my baseline estimate for the same overlapping sample, I find quantitatively very close hour-adjusted supply elasticity estimates across the productivity distribution. Thus the hours margin of adjustment also does not appear to play a major quantitative role in practice.

I turn next to the observable firm-specific labor supply shocks introduced in section [2.3](#). In appendix Table [IA.11](#) I show that controlling for two of the most prominent examples of such supply shocks identified in the literature—unionization events or changes in non-compete enforceability—has essentially no effect on the average supply elasticity estimate. This is in accordance with the theme across all these robustness exercises: my baseline average supply elasticity estimate is always very close economically and statistically to the various alternatives, suggesting that any bias in my estimates is likely to be quantitatively small.

I also explore numerous robustness checks for labor productivity-sorted supply elasticity estimates; for further details on all these estimates see internet appendix section [IA.7](#). In appendix Table [IA.8](#) I sort firms into labor productivity quartiles within 2-digit NAICS industry. In appendix Table [IA.9](#) I perform several more robustness checks on elasticities sorted on productivity. I first estimate elasticities at the 3-year instead of 1-year horizon to test whether short term frictions drive the lower elasticity estimates for unproductive firms. Next, I replace average log wage changes with time-varying AKM firm wage effects to verify that sorting patterns are not likely to be driven by differences in incentive pay or match components of wages. Then I sort firms on TFP estimates from [İmrohoroglu and Tüzel \(2014\)](#) instead of value added per worker. Finally, to address concerns that Compustat employment may reflect workers in non-domestic establishments, I replace Compustat employment changes with employment changes from the Census Longitudinal Business Database (LBD). In all cases I find the same decreasing pattern in supply elasticities as productivity increases, while my use of Compustat employment turns out to be a quantitatively conservative choice.

Besides these robustness checks, another concern could be that the use of stock returns as a shock to labor demand systematically biases me to find these cross-sectional sorting patterns, especially if more productive firms tend to compensate workers with stock-based compensation for long-term contracting reasons unrelated to a wage posting, monopsonistic framework.²³ However, it is actually smaller firms—which are less productive on average—that tend to offer relatively more stock-based compensation (see [Eisfeldt, Falato, and Xiaolan \(2021\)](#) for evidence on this point).

²³Note that stock-based pay is not necessarily inconsistent with a monopsony model, as it naturally increases compensation when firm performance—and hence marginal productivity—is higher, exactly as implied in monopsony models.

Two additional tests confirm that differences in stock-based incentive pay are not driving my cross-sectional findings. In appendix Table [IA.10](#) I show that sorting patterns in wage responses to stock return shocks are similar for both positive and negative shocks; in appendix Figure [IA.6](#) I follow [Eisfeldt et al. \(2021\)](#) to create a proxy for ex-ante intensity of reliance on stock-based compensation, incentive pay, or perks for skilled employees, and I find nearly the exact same estimates when sorting on labor productivity within groups of firms that are high- and low-intensity for these types of compensation. For further details on these estimates and the rationale behind them, refer to internet appendix section [IA.7.2](#).

3.4.1 Other Explanations for My Findings

[Hartman-Glaser et al. \(2019\)](#) and [Kehrig and Vincent \(2021\)](#) examine related patterns in firm-level labor shares (namely, a widening cross-sectional spread that contributes to the aggregate time series decline) but argue for different economic mechanisms. The former contend that increased idiosyncratic risk has played a crucial role, while the latter argue their findings are consistent with reallocations of output to firms with large price markups. In internet appendix section [IA.10.1](#) I explore why the idiosyncratic risk mechanism from [Hartman-Glaser et al. \(2019\)](#) can't account for labor share patterns within my sample period, and why my empirical estimates are not inconsistent with those of [Kehrig and Vincent \(2021\)](#).

3.5 Summary of empirical findings

Overall, the estimates in this section imply that labor market power is likely to be a quantitatively important determinant of firm cashflows. In particular, rents from labor market power benefit capital owners in the form of higher valuations and lower labor shares, especially for the most productive firms. Skilled workers appear likely to play an especially important quantitative role, as skilled workers at productive firms exhibit the least elastic labor supply of all.

While my estimates are qualitatively informative about the magnitude of labor market power, they may not be sufficient by themselves to quantify the cashflows that firms derive from wage markdowns. As discussed before, in the presence of labor adjustment costs the supply elasticity is not the only determinant of markdowns from marginal product. I now formalize this concept in section [4](#), where I show how to map empirical supply elasticity estimates into quantitative markdowns in the presence of labor market adjustment costs, by exploiting firm Euler equations over empirically observable outcomes.

4 Wage Markdowns Under Dynamic Monopsony

In this section I map out a dynamic extension to the static framework described in section [2](#), which will allow me to calculate dynamically-adjusted wage markdowns. Consider the

following dynamic modification of the firm's problem (1):

$$V_t(L_t; \Theta_t, \Gamma_t) = E_t \left[\sum_{t=0}^{\infty} \rho^t \left(\max_{L_t} F_t(L_t; \Theta_t) - W_t(L_t; \Gamma_t) L_t - \Phi_t(L_t, L_{t-1}) \right) \right] \quad (17)$$

The function $\Phi_t(L_t, L_{t-1})$ represents adjustment costs of current labor, which depends on the firm's labor force in the previous period, making the firm's optimal labor choice dynamic. Likewise, the inverse labor supply curve $W_t(L_t, L_{t-1}; \Gamma_t)$ depends on last period's labor. To see how this changes the firm's optimal wage markdown, I take the first-order condition to get the Euler equation:

$$\frac{\partial F_t}{\partial L_t} - \frac{\partial \Phi_t}{\partial L_t} = \frac{\partial W_t}{\partial L_t} L_t + W_t + \rho E_t \left[\frac{\partial \Phi_{t+1}}{\partial L_t} \right] \quad (18)$$

Divide both sides by W_t , and define the net inverse wage markdown $\Pi_t \equiv \frac{1}{W_t} \left(\frac{\partial F_t}{\partial L_t} - \frac{\partial \Phi_t}{\partial L_t} \right)$. This yields the modified Euler equation

$$\Pi_t = \frac{1 + \epsilon_{L_t, W_t}}{\epsilon_{L_t, W_t}} + \frac{1}{W_t} \rho E_t \left[\frac{\partial \Phi_{t+1}}{\partial L_t} \right], \quad (19)$$

The net inverse wage markdown Π_t is the marginal product of labor (net of labor overhead costs) divided by the wage. When Π_t is larger, the firm generates more income from labor market power. In accordance with the usual static notion, Π_t is decreasing in ϵ_{L_t, W_t} . Importantly, if Φ_t is strictly decreasing in last period's labor choice L_{t-1} , then the last term on the right-hand side of (19) is negative, and profits from wage markdowns are strictly smaller under (19) than in the static setting. In the case where the term in the expectation is zero, the cost function Φ_t is static, and (19) collapses to the usual first-order condition (2).

4.1 Dynamic monopsony: An illustrative example

Does dynamic monopsony seem consistent with the data? In internet appendix section [IA.1.1](#), I map out a straightforward dynamic extension of the example monopsony model from section [2.2.1](#). The model is a simple Q-theory framework for both labor and capital. The labor supply functional form again borrows from [Kline et al. \(2019\)](#) and [Card et al. \(2018\)](#), and I incorporate quadratic labor hiring and capital adjustment costs in a setting where the firm gradually accumulates labor. See internet appendix [IA.1.1](#) for all details on the setup, first-order conditions, and calibration.

I simulate the impact of a small one-time, mean-reverting TFP shock, where I calibrate the TFP shock persistence to exactly match the decay in the empirical response to stock returns in [Figure 1](#). I further calibrate the model to match the pooled, homogeneous labor supply elasticity of about 2.5 from column (1) of [Table 1](#); and the dynamic growth responses of employment, wages, and capital observed in [Figure 1](#) and appendix [Figure IA.2](#) over the

next five years. As can be seen in appendix Figure [IA.3](#), the model matches the empirical dynamic responses quite well. The model further generates a worker separation rate, firm log labor share, and magnitudes of initial wage and employment responses (relative to the size of the initial TFP shock) that are closely in line with their empirical counterparts. Finally, the model implies lower supply elasticities for high productivity firms, and higher supply elasticities for low productivity firms. The model-implied spread in supply elasticities resembles the empirical spread when the model TFP gap is calibrated to roughly align with the empirical TFP gap between above- and below-median labor productivity firms. See appendix section [IA.1.1](#) for elaboration on the above points.

Still, this simple deterministic steady-state partial equilibrium setup is too stylized to make quantitative predictions about wage markdowns, and it would be difficult to estimate when adding more dimensions of worker heterogeneity. Instead, motivated by the above evidence suggesting that a simple dynamic monopsony model with quadratic hiring costs provides a good approximation to the data, I now directly estimate the firm’s dynamic Euler equation (19) by exploiting a few informative empirical moments and the same adjustment cost functional form.

4.2 Mapping dynamic framework to the data

How much do dynamic adjustments matter quantitatively? We can make progress by noticing that most of the Euler equation (19) can be mapped to the data. We have already addressed estimation of ϵ_{L_t, W_t} in detail. This leaves ρ and the adjustment cost function to be identified. The parameter ρ corresponds to the real discount rate at the one-year horizon, and should naturally take a value that’s close to 1. For this reason, my quantitative findings are not sensitive to any reasonable choice for ρ . I choose $\rho = \exp(-0.065)$, corresponding to an average nominal discount rate of approximately 9% (roughly the average rate reported in [\(Decaire and Graham, 2024\)](#)), minus an average inflation rate of about 2.5% over my sample period.

The adjustment cost function $\Phi(L_t, L_{t-1})$ is the last ingredient required to estimate the firm’s Euler equation. As in [\(Belo, Gala, Salomao, and Vitorino, 2022; Acemoglu and Hawkins, 2014\)](#), and consistent with the calibration of the dynamic model described in the previous section, I assume adjustment costs are a quadratic function of the firm’s gross hiring rate.²⁴ In internet appendix section [IA.1.3](#), I show that an extended version of the dynamic monopsony framework—where firms can now make separate wage offers to incumbent workers and new hires—contains an additional moment which allows me to estimate the last adjustment cost parameter γ_L , which maps the gross hiring rate into marginal adjustment costs. Accordingly,

²⁴Recent work offers direct evidence that firms find it costly to replace incumbent workers with new hires [\(Kline et al., 2019; Jäger et al., 2024\)](#); I offer some supporting evidence from my data in internet appendix section [IA.1.4](#).

I exploit this moment to recover $\frac{1}{W_t} E_t \left[\frac{\partial \Phi_{t+1}}{\partial L_t} \right]$ in (19).²⁵ See appendix [IA.1.4](#) for all estimation details.

An advantage of this approach is that I can estimate (19) flexibly across worker and firm types and over time. In the next section I explore the quantitative implications of applying (19) with heterogeneity by firm productivity, time period, and worker skill.

5 Quantifying Cashflows Generated from Labor Market Power

In this section I use the estimated parameters of (19) to quantify the income that firm capital holders generate from labor market power. I estimate (19) separately by firm labor productivity (above- versus below-median) interacted with time period (1991-2002 versus 2003-2014). I explore both a homogeneous worker setup, which I take as my baseline, and a version where I estimate (19) with worker skill heterogeneity, following the same skill breakdown used in my empirical estimates in Tables 1 and 3.

5.1 Adjustment costs

Before analyzing the value of wage markdowns, it is instructive to look at the implied average marginal adjustment costs—given by the product of my estimates of γ_L and firms’ hiring rates—implied by my procedure. As discussed in online appendix [IA.1.4](#), γ_L can be calibrated by estimating the relationship between the (supply elasticity-adjusted) wage premium of incumbent employees over new hires with the firm’s hiring rate.

I report pooled means and medians of the implied adjustment costs in appendix Table [IA.6](#). In the first column I show the estimated real marginal adjustment cost per new hire, which I compare to average wages in the next two columns. The first panel reports estimates for the version with homogeneous worker types (though I do still allow for γ_L to vary by firm productivity and time period). Examining the top row, the full-sample mean marginal adjustment cost is about \$40,000 (expressed in real 2011 US dollars), which is 58% of the average wage, while median adjustment costs are about 47% of the median wage. This average masks heterogeneity across time and firm type. Adjustment costs have risen over time, from just under half to about 70% of the typical wages when comparing the first and second half of the sample. Meanwhile, productive firms experience slightly higher marginal adjustment costs relative to wages overall. While these magnitudes are economically substantial, they are also consistent with other work: [Kline et al. \(2019\)](#) find adjustment costs that are just

²⁵I just use this single moment from the extended version with incumbent and recruit heterogeneity to inform my calibration of γ . A complete quantification of the extended version is quite sensitive to the gap in the supply elasticities estimated separately for incumbents and recruits combined with the composition of incumbents and recruits across firms and over time. Still, it turns out that estimating a fully-extended version with separate offers to incumbents and recruits, plus worker skill heterogeneity, implies that firms generate a similar, but slightly larger, share of profits from labor market power as my baseline—albeit with a much larger time-series increase between the first and second halves of the sample.

above the average wage of new hires, and Jäger et al. (2024) estimate that adjustment costs are about two-thirds to more than twice a replaced worker’s wages, though both focus on samples with mostly smaller firms.

In panel B, I allow for heterogeneity in γ_L based on worker skill, and a stark pattern emerges. Adjustment costs clearly increase in worker skill: low- and middle-wage workers’ adjustment costs are only 3 to 8% of wages, respectively, while high-skilled workers’ adjustment costs are about one and half times the average worker wage (or roughly equal to wages at the median). This makes economic sense to the extent that it is far more costly to recruit, train, and integrate workers with highly-specialized skills into the production process. These estimates in turn imply that dynamic adjustments to wage markdowns in (19) apply primarily to the high-skill workers, while labor determination is likely to be nearly static for low- and middle-skill employees. I now examine what the dynamic adjustments imply for firm profits by applying (19) to calculate the cash flows generated from labor market power.

5.2 Valuing Labor Market Power: Quantitative Implications for Firm Cash Flows, Labor Shares, and Valuations

I now examine how firm profits from labor market power affect aggregate and cross-firm outcomes. This forges a direct quantitative link between estimated wage markdowns and the large gaps in profits, labor shares, and valuations across the productivity distribution. I also examine the role of labor market power in explaining the time series change in the labor share.

5.2.1 Labor market power profits as a share of capital income

In Table 6 I explore the quantitative implications of my findings for firm cash flows. I begin by focusing on the first two columns, which report pooled firm-level average profits per worker and net inverse wage markdowns for different sets of firms and time periods. The first column of Table 6 computes average real profits per worker (in 2011 US dollars), where profits per worker are given by the empirical estimates of $(\Pi_t - 1) \times W_t$.²⁶ Focusing first on panel A, the full-sample average profits per employee is around \$17,550, which rose from \$10,570 in the first half of the sample to \$25,840 in the second half. Cross-sectionally, productive firms generate substantially higher profits per worker (\$4,891 for below median productivity versus \$30,210 for above median). In the first column of panel B, I allow for skill heterogeneity in markdowns. I find considerable heterogeneity across skill levels: profits

²⁶Profits per worker are equal to the gap between net marginal revenues and the wage:

$$\text{Per worker profits} = \frac{\partial F}{\partial L} - \frac{\partial \Phi}{\partial L} - W_t = \Pi_t \times W_t - W_t = (\Pi_t - 1) \times W_t \quad (20)$$

Rather than taking the firm-level average full-time, full-year equivalent adjusted wages that were used to estimate supply elasticities, here I take W_t to be the raw LEHD wage bill per worker. This measure of wages is more appropriate in this context, as it reflects actual average spending per employee.

per worker are respectively from \$3,923, \$8,678, and \$39,910 for low-, middle, and high-skill workers, respectively. Comparing the amount for high skill workers to their average wage (\$142,300, reported in Table IA.6), this implies that the typical skilled worker contributes about \$180,000 of net revenues on the margin, of which the average firm takes nearly \$40,000 as profits.

In the second column of Table 6 I report the pooled averages of Π_t , while the third column reports the supply elasticity that would be consisting with this markdown in a static model. The average of Π_t across all firm-years is about 1.22. The first row of the third column shows that this is consistent with a supply elasticity of about 4.5, which is considerably higher than the empirical homogeneous elasticity of about 2.5 (reported again in the fourth column for ease of comparison). Quite interestingly, this implied elasticity of 4.5 is nearly the exact same as the elasticity of 4.6 implied by the calibration of the dynamic monopsony model in appendix IA.1.2. Implied average net inverse markdowns rise over time (1.16 in 1991-2002 to 1.30 in 2003-2014) and increase with productivity (1.09 for low productivity and 1.35 for high). In panel B I explore average estimates of Π_t with worker skill heterogeneity. After averaging within the firm across skill types to get implied firm profits per worker under skill heterogeneity, the implied average firm profits per worker are similar at a little over \$17,000. The overall average markdown is now 1.20 (computed from the wagebill-weighted average of Π_t across skill groups within the firm). Separating out by skill, pooled average markdowns rise from 1.14 for low-skill workers to 1.23 for high-skill workers. While the slope in markdowns by worker skill is sizable, it's smaller than would be implied by supply elasticity estimates alone. In fact, skilled workers face a markdown that's consistent with a supply elasticity estimate of above 4, despite an empirical supply elasticity of just above 1. Meanwhile, the elasticity implied by the low- and middle-skill workers' markdowns is far closer to the empirical estimate. This underscores the importance of adjustment costs dynamics for skilled workers, where ignoring dynamics implies a large overestimate of market power. By contrast, a static markdown seems a reasonable approximation for low- and middle-skill employees.

In the final three columns of Table 6, I compare the profits earned from market power to aggregate capital income and revenues. My baseline estimate of the average yearly value of labor market profits divided by total capital income (where capital income is defined as operating income before depreciation plus changes in inventories) is about 25%, which can be found in the first row of the fifth column. This baseline measure of capital income may neglect some intangible investments that are expensed rather than treated as capital investments. In the next column I report the share of labor market profits relative to an adjusted measure of capital income which adds back in only the imputed non-labor component of intangible investments. This is important because the labor share of intangible investments is quite

high.²⁷ Labor market profits are about 22% of adjusted capital income. In the last column, I compare markdown profits to firm revenues, finding they constitute close to 4% of aggregate yearly revenues on average.

Again, these figures mask heterogeneity in the time series and the cross section of firms. The share of capital income earned from labor market profits rose from 20% in the first half of the sample to roughly 31% in the second half (from 17% to 28% when using adjusted capital income, and from about 3% to 5% of revenues). Meanwhile, productive firms’ wage markdowns represent 28% of capital income, compared to 13% for unproductive firms. Moving to analyze skill heterogeneity in panel B, I find that wage markdowns are worth a similar implied share of aggregate capital income at 23% (or 20% for adjusted capital income); what’s more, skilled workers account for nearly 70% of the value generated from this channel, underscoring the importance of highly-skilled employees.

5.2.2 Labor market power and labor shares

Next, in Table 7 I examine how labor market power has affected labor shares in the cross-section (panel A) and time series (panel B). Theoretically, the firm-level log labor shares should move one-for-one with the log wage markdown. In panel A, I leverage this fact to quantify how much markdown differences can explain average log labor share differences for high- versus low-productivity firms. In the first row of panel A, I report the average log labor share of low- and high-productivity firms, which are -0.30 and -0.73, respectively, with an economically very large gap of 0.43 log points. In the next row I take the log of one over the average net inverse markdown estimated for the two productivity types; these are -0.09 and -0.30, respectively, with a gap of about 0.21 log points. This in turn implies that the difference in log wage markdowns can explain $0.21/0.43 \approx 50\%$ of the wide cross-sectional labor share gap. The share of the cross-sectional gap explained under the skill heterogeneity counterfactual is slightly smaller but still substantial, at about 35%. Thus labor market power variation is indeed a quantitatively important determinant of this large cross-sectional firm labor share spread.

In panel B of Table 7, I examine the role of labor market power in explaining the time series labor share decline. I calculate that the public firm labor share declined from 57.9% in 1991-2002 to 52.4% in 2003-2014, which I show in the first row of panel B. In the next rows I consider the following counterfactual: “What would the labor share have been in the 2003-2014 period if we re-calculate implied wage markdowns in (19), assuming labor supply elasticities and adjustment cost parameters were held constant at their estimated values in 1991-2002?” This is a partial equilibrium exercise because it still holds firm employment

²⁷Lehr (2025) calculates that about 70% of US R&D spending in the year 2000 was on salary and wages paid to employees involved in the R&D process. Using this proxy for the labor share in intangible investments, I add back in 30% of expensed intangible investment—defined as 30% of SG&A plus R&D following the literature (Eisfeldt and Papanikolaou, 2013; Peters and Taylor, 2017; Crouzet and Eberly, 2021)—to arrive at my measure of adjusted capital income.

decisions constant.²⁸ For the homogeneous worker version, I find the average aggregate labor share would have declined to about 54.4% in 2003-2014, implying that my estimates explain about 35% of the time series decline. While the skill heterogeneity counterfactual explains a smaller portion of the cross-sectional labor share spread, the opposite is true of the aggregate time series: with worker skill heterogeneity I find the aggregate labor share would have declined to only 55.8% in 2003-2014, explaining about three-fifths of the time series change.

5.2.3 Labor market power, firm profit rates, and valuations

In Table 2 I showed that productive firms have high valuations and profits, but they do not invest more; Table 5 established a qualitative association between elasticity-implied static markdowns and these same firm outcomes. In Table 8, I complete the connection by examining the quantitative role of labor market power for profits and valuations. In the top panel of Table 8, I look at firm profitability, as defined by firms' return on assets. I compute the counterfactual return on assets when I remove firms' labor market power profits from their income, which I then compare to their actual return on assets. The first column of panel A shows that low productivity firms have an average return on assets of 9.5%, while productive firms have a return on assets of 16.5%, with a gap of 7 percentage points. In the second column, I find counterfactual return on assets of 7% and 10% for unproductive and productive firms, respectively, after removing labor market power cash flows. This closes the profitability gap to almost 3 percentage points, implying the labor market income to capital owners can explain 58% of the average gap in this profitability measure. When I use the skill heterogeneity counterfactual, I explain a higher share of unproductive firms' profit rates (their profits now decline from 9.5% to 6.6%), which implies that I close a smaller share of the cross-sectional gap at 36%.

In panel B of the table, I examine the impact on firm valuations, again using the [Peters and Taylor \(2017\)](#) Tobin's Q ratio that includes intangibles. To do this exercise, I follow [Crouzet and Eberly \(2021\)](#), who show that Tobin's Q can be decomposed into an economic rents component driven by firm profitability (return on assets) as well as the true marginal investment incentive. I implement an empirical analogue of their decomposition by estimating cross-sectional regressions of firms' valuation ratios on return on assets and actual firm-level investment rates. I then use estimated regression coefficients to predict the counterfactual firm valuation ratio when I hold investment rates constant, but remove labor market power profits from firms' return on assets. See internet appendix section [IA.2](#) for full details on the estimation procedure. Importantly, because I control for investment rates, the counterfactual change in valuations is driven by the economic rents component of valuations, and not the component of valuations driven by investment incentives.

In the first column of panel B, I find an average valuation of 0.77 for unproductive

²⁸I cannot make statements about the counterfactual labor share in general equilibrium because I don't have a full equilibrium structural model estimate.

firms and 1.38 for productive firms. After removing labor market power cashflows, these valuations decline to 0.67 and 0.95 for unproductive and productive firms, respectively, in the homogeneous worker decomposition. This implies that labor market power can explain 54% of the cross-sectional gap in valuation ratios. Under worker skill heterogeneity, I explain almost two-fifths of the gap; the slightly smaller share explained is again because I estimate larger labor market profits for unproductive firms under skill heterogeneity. As highlighted in Table 2, productive firms’ high valuations are not met by high investment rates. By explaining much of these average difference in valuation ratios, labor market power is a quantitatively important determinant of this cross-sectional “investment gap”.

5.3 Discussion

This section has quantified the role of labor market power in driving cross-sectional differences in average valuations, labor shares, and profit rates, as well as the time series decline in the labor share of income, among publicly traded firms. Despite the fact that dynamic considerations temper the exercise of market power—especially for skilled workers—I find a quantitatively meaningful role for labor market power in explaining these differences.

Given the quantitative magnitude of my findings, it’s worth benchmarking against estimates in the broader literature on market power. I do this in section IA.9 of the internet appendix. As I discuss there, my dynamically-adjusted estimates of wage markdowns are on the lower end of markdown or markup estimates. The reason labor market power is quantitatively important for firm profits is not because my markdown estimates are large. Rather, it is because labor is the largest expense for most firms, such that a modest amount of market power over labor can generate considerable value. This is why it can explain a large chunk of wide empirical gaps in profits, valuations, and labor shares between productive and unproductive firms. Finally, I stress that my findings still leave plenty of room for other economic forces, such as increasing price markups or technological change, to explain cross-sectional differences and time series changes in labor shares. I do argue, however, that labor market power has played a more important role in these phenomena than prior literature may suggest.

6 Conclusion

In this paper I find evidence that “superstar firms” have generated substantial and increasing value from their labor market power in the United States. Firms with productivity advantages face much lower supply elasticities and hence earn higher monopsony rents from wage markdowns. Differences in markdowns have widened over time, leading to an increased gap in labor shares between productive and unproductive firms in the cross section, and contributing to the decline in the aggregate labor share. The overall value of wage markdowns is substantial. They are worth about 25% of aggregate capital income to publicly traded

firms under homogeneous worker types and 23% once I allow markdowns to differ by worker skill (22% and 20% respectively against intangible-adjusted capital income), or roughly 4% of aggregate revenues; the figure is 28% among the more productive half of firms and 13% among the less productive half. Removing these cash flows closes 58% of the cross-sectional gap in operating profitability and 54% of the gap in valuation ratios between productive and unproductive firms, and 36% and roughly two-fifths of those gaps respectively when I allow for skill heterogeneity.

Two elements of the approach may be useful beyond this application. The first is the use of idiosyncratic stock returns as a labor demand shifter. Returns are available for every public firm at high frequency and impound a broad range of news about a firm’s prospects, which makes them a more powerful instrument for labor demand than any single observable shock, and the identification argument in section 2 makes clear what they do and do not require. The second is the correction for labor adjustment costs. A static supply elasticity is not sufficient to recover a firm’s wage markdown when labor is costly to adjust, and the dynamic term turns out to be first-order precisely for the workers who generate most of the rents. Ignoring it would have overstated skilled workers’ markdowns considerably.

Skilled workers are the key quantitative driver of these findings, which points to the need to examine the frictions in the high-skill labor market, since they may be very different than for other workers. Labor markets are much less local for skilled workers (see [Malamud and Wozniak, 2012](#); [Amior, 2020](#), for example), which limits how much concentration measured over narrow geographic boundaries can say about them. Firm-specific human capital may instead play an outsize role. Supporting a human capital specificity mechanism, [Nimczik \(2020\)](#) finds that highly-skilled workers tend to be employed in labor markets that are geographically dispersed but more concentrated in a set of specific industries. The increasing importance of firm-specific human capital could also be a common explanation behind both the rising adjustment costs and the declining supply elasticities that I find. My evidence on this front is only indirect, and pinning down the specific frictions that skilled workers face is the natural next step: it is where the large majority of the rents documented here is generated, and it is the part of the labor market that existing concentration-based measures are least equipped to describe.

References

- Abel, A. B. and J. C. Eberly (2011, 03). How Q and Cash Flow Affect Investment without Frictions: An Analytic Explanation. *The Review of Economic Studies* 78(4), 1179–1200.
- Abowd, J. A. and T. Lemieux (1993). The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada. *The Quarterly Journal of Economics* 108(4), 983–1014.
- Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica* 67(2), 251–333.
- Abowd, J. M., B. E. Stephens, L. Vilhuber, F. Andersson, K. L. McKinney, M. Roemer, and

- S. Woodcock (2009). *The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators*. University of Chicago Press.
- Acemoglu, D. and W. B. Hawkins (2014). Search with multi-worker firms. *Theoretical Economics* 9(3), 583–628.
- Amior, M. (2020). Education and geographical mobility: The role of the job surplus.
- Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020, 02). The Fall of the Labor Share and the Rise of Superstar Firms. *The Quarterly Journal of Economics* 135(2), 645–709.
- Azar, J., I. Marinescu, and M. Steinbaum (2022). Labor market concentration. *Journal of Human Resources* 57(S), S167–S199.
- Bachmann, R., G. Demir, and H. Frings (2022). Labor market polarization, job tasks, and monopsony power. *Journal of Human Resources* 57(S), S11–S49.
- Balke, N. and T. Lamadon (2020, November). Productivity shocks, long-term contracts and earnings dynamics. Working Paper 28060, National Bureau of Economic Research.
- Barkai, S. (2020). Declining labor and capital shares. *The Journal of Finance* 75(5), 2421–2463.
- Bassier, I., A. Dube, and S. Naidu (2022). Monopsony in movers: The elasticity of labor supply to firm wage policies. *Journal of Human Resources* 57(S), S50–S86.
- Belo, F., V. D. Gala, J. Salomao, and M. A. Vitorino (2022). Decomposing firm value. *Journal of Financial Economics* 143(2), 619–639.
- Belo, F., X. Lin, and S. Bazdresch (2014). Labor hiring, investment, and stock return predictability in the cross section. *Journal of Political Economy* 122(1), 129–177.
- Benmelech, E., N. Bergman, and H. Kim (2022). Strong employers and weak employees: How does employer concentration affect wages? *Journal of Human Resources* 57(S), S200–S250.
- Berger, D. W., K. F. Herkenhoff, and S. Mongey (2022). Labor market power. *American Economic Review* 112(4), 1147–1193.
- Bradsen, G. and H. Lütkepohl (2011). Forecasting levels of log variables in vector autoregressions. *International Journal of Forecasting* 27(4), 1108–1115.
- Caldwell, S. and E. Oehlsen (2018). Monopsony and the gender wage gap: Experimental evidence from the gig economy.
- Card, D. (2022). Who set your wage? *American Economic Review* 112(4), 1075–1090.
- Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics* 36(S1), S13–S70.
- Card, D., J. Heining, and P. Kline (2013, 05). Workplace Heterogeneity and the Rise of West German Wage Inequality. *The Quarterly Journal of Economics* 128(3), 967–1015.
- Chan, L. K. C., N. Jegadeesh, and J. Lakonishok (1996). Momentum strategies. *The Journal of Finance* 51(5), 1681–1713.
- Chan, M., S. Salgado, and M. Xu (2021). Heterogeneous passthrough from tfp to wages.
- Cho, D. (2018). The labor market effects of demand shocks: Firm-level evidence from the recovery act.
- Cohen, L. and A. Frazzini (2008). Economic links and predictable returns. *The Journal of Finance* 63(4), 1977–2011.
- Corhay, A., H. Kung, and L. Schmid (2020). Q: Risk, rents, or growth?
- Covarrubias, M., G. Gutiérrez, and T. Philippon (2020). From good to bad concentration? us industries over the past 30 years. *NBER Macroeconomics Annual* 34, 1–46.
- Crouzet, N. and J. Eberly (2021). Rents and intangible capital: A q+ framework.
- Daniel, K., D. Hirshleifer, and L. Sun (2019, 06). Short- and Long-Horizon Behavioral Factors. *The Review of Financial Studies* 33(4), 1673–1736.
- Davis, S. J., J. Haltiwanger, R. Jarmin, J. Miranda, C. Foote, and E. Nagypal (2006). Volatility and dispersion in business growth rates: Publicly traded versus privately held firms [with comments


- and discussion]. *NBER Macroeconomics Annual* 21, 107–179.
- De Loecker, J., J. Eeckhout, and G. Unger (2020, 01). The Rise of Market Power and the Macroeconomic Implications. *The Quarterly Journal of Economics* 135(2), 561–644.
- Decaïre, P. and J. Graham (2024). Valuation fundamentals.
- Donangelo, A. (2014). Labor mobility: Implications for asset pricing. *The Journal of Finance* 69(3), 1321–1346.
- Donangelo, A., F. Gourio, M. Kehrig, and M. Palacios (2019). The cross-section of labor leverage and equity returns. *Journal of Financial Economics* 132(2), 497–518.
- Dube, A., A. Manning, and S. Naidu (2018, September). Monopsony and employer mis-optimization explain why wages bunch at round numbers. Working Paper 24991, National Bureau of Economic Research.
- Eisfeldt, A. L., A. Falato, and M. Z. Xiaolan (2021, May). Human capitalists. Working Paper 28815, National Bureau of Economic Research.
- Eisfeldt, A. L. and D. Papanikolaou (2013). Organization capital and the cross-section of expected returns. *The Journal of Finance* 68(4), 1365–1406.
- Engbom, N. and C. Moser (2020). Firm pay dynamics.
- Ewens, M. and M. Marx (2017). Founder replacement and startup performance. *The Review of Financial Studies* 31(4), 1532–1565.
- Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics* 116(1), 1–22.
- Farhi, E. and F. Gourio (2018, November). Accounting for macro-finance trends: Market power, intangibles, and risk premia. Working Paper 25282, National Bureau of Economic Research.
- Favilukis, J. and X. Lin (2015, 08). Wage rigidity: A quantitative solution to several asset pricing puzzles. *The Review of Financial Studies* 29(1), 148–192.
- Favilukis, J., X. Lin, and X. Zhao (2020, June). The elephant in the room: The impact of labor obligations on credit markets. *American Economic Review* 110(6), 1673–1712.
- Ferres, D., G. Kankanalli, and P. Muthukrishnan (2023). Anti-poaching agreements, innovation, and corporate value: Evidence from the technology industry.
- Friedrich, B., L. Laun, C. Meghir, and L. Pistaferri (2019, April). Earnings dynamics and firm-level shocks. Working Paper 25786, National Bureau of Economic Research.
- Garin, A. and F. Silverio (2020). How responsive are wages to demand within the firm? evidence from idiosyncratic export demand shocks.
- Gouin-Bonenfant, É. (2022). Productivity dispersion, between-firm competition, and the labor share. *Econometrica* 90(6), 2755–2793.
- Greenwald, D., M. Lettau, and S. Ludvigson (2021). How the wealth was won: Factor shares as market fundamentals.
- Grullon, G., Y. Larkin, and R. Michaely (2019, 04). Are US Industries Becoming More Concentrated?*. *Review of Finance* 23(4), 697–743.
- Guiso, L., L. Pistaferri, and F. Schivardi (2005). Insurance within the firm. *Journal of Political Economy* 113(5), 1054–1087.
- Gutierrez, G. and T. Philippon (2017). Investmentless growth: An empirical investigation. *Brookings Papers on Economic Activity*, 89–169.
- Hagedorn, M. and I. Manovskii (2008, September). The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review* 98(4), 1692–1706.
- Hartman-Glaser, B., H. Lustig, and M. Z. Xiaolan (2019). Capital share dynamics when firms insure workers. *The Journal of Finance* 74(4), 1707–1751.
- Howell, S. T. and J. D. Brown (2020, January). Do cash windfalls affect wages? evidence from r&d grants to small firms. Working Paper 26717, National Bureau of Economic Research.
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in china and india. *The*

- Quarterly Journal of Economics* 124(4), 1403–1448.
- Jäger, S., J. Heining, and N. Lazarus (2024). How substitutable are workers? evidence from worker deaths. Working Paper 30629, National Bureau of Economic Research. Conditionally accepted, *American Economic Review*.
- Jarosch, G., J. S. Nimczik, and I. Sorkin (2024). Granular search, market structure, and wages. *The Review of Economic Studies* 91(6), 3569–3607.
- Jeffers, J. (2019). The impact of restricting labor mobility on corporate investment and entrepreneurship.
- Karabarbounis, L. and B. Neiman (2013, 10). The Global Decline of the Labor Share*. *The Quarterly Journal of Economics* 129(1), 61–103.
- Kehrig, M. and N. Vincent (2021, 03). The Micro-Level Anatomy of the Labor Share Decline. *The Quarterly Journal of Economics* 136(2), 1031–1087.
- Kim, H. (2020). How does labor market size affect firm capital structure? evidence from large plant openings. *Journal of Financial Economics* 138(1), 277–294.
- Kline, P., N. Petkova, H. Williams, and O. Zidar (2019, 03). Who Profits from Patents? Rent-Sharing at Innovative Firms. *The Quarterly Journal of Economics* 134(3), 1343–1404.
- Knepper, M. (2020, 03). From the fringe to the fore: Labor unions and employee compensation. *The Review of Economics and Statistics* 102(1), 98–112.
- Kogan, L., D. Papanikolaou, L. D. W. Schmidt, and J. Song (2020, April). Technological innovation and labor income risk. Working Paper 26964, National Bureau of Economic Research.
- Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017, 03). Technological Innovation, Resource Allocation, and Growth*. *The Quarterly Journal of Economics* 132(2), 665–712.
- Kroft, K., Y. Luo, M. Mogstad, and B. Setzler (2025). Imperfect competition and rents in labor and product markets: The case of the construction industry. *American Economic Review* 115(9), 2926–2969.
- Kuehn, L.-A., M. Simutin, and J. J. Wang (2017). A labor capital asset pricing model. *The Journal of Finance* 72(5), 2131–2178.
- Lachowska, M., A. Mas, R. D. Saggio, and S. A. Woodbury (2020, January). Do firm effects drift? evidence from washington administrative data. Working Paper 26653, National Bureau of Economic Research.
- Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review* 112(1), 169–212.
- Lee, D. S. and A. Mas (2012, 01). Long-run impacts of unions on firms: New evidence from financial markets, 1961–1999. *The Quarterly Journal of Economics* 127(1), 333–378.
- Lehr, N. (2025). Does monopsony matter for innovation?
- Lindenlaub, I., R. Oh, and M. Peters (2024). Spatial firm sorting and local monopsony power.
- Lipsius, B. (2018). Labor market concentration does not explain the falling labor share.
- Liu, Y. (2019). Labor based asset pricing.
- Malamud, O. and A. Wozniak (2012). The impact of college on migration: Evidence from the vietnam generation. *The Journal of Human Resources* 47(4), 913–950.
- Matsa, D. A. (2010). Capital structure as a strategic variable: Evidence from collective bargaining. *The Journal of Finance* 65(3), 1197–1232.
- McFadden, D. L. (1973). Conditional logit analysis of qualitative choice behavior. In *Frontiers in econometrics*, pp. 105–42. Academic Press: New York.
- Mertens, M. (2023). Labor market power and between-firm wage (in)equality. *International Journal of Industrial Organization* 91, 103005.
- Mueller, H. M., P. P. Ouimet, and E. Simintzi (2017, 05). Within-Firm Pay Inequality. *The Review of Financial Studies* 30(10), 3605–3635.
- Nimczik, J. (2020). Job mobility networks and data-driven labor markets.

- Pellegrino, B. (2025, April). Product differentiation and oligopoly: A network approach. *American Economic Review* 115(4), 1170–1225.
- Peters, M. (2020). Heterogeneous markups, growth, and endogenous misallocation. *Econometrica* 88(5), 2037–2073.
- Peters, R. H. and L. A. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics* 123(2), 251–272.
- Ransom, M. R. and D. P. Sims (2010). Estimating the firm’s labor supply curve in a “new monopsony” framework: Schoolteachers in missouri. *Journal of Labor Economics* 28(2), 331–355.
- Rinz, K. (2022). Labor market concentration, earnings, and inequality. *Journal of Human Resources* 57(S), S251–S283.
- Robinson, J. (1969). *The Economics of Imperfect Competition*. Macmillan.
- Schubert, G., A. Stansbury, and B. Taska (2024). Employer concentration and outside options.
- Shen, M. (2021, 02). Skilled Labor Mobility and Firm Value: Evidence from Green Card Allocations. *The Review of Financial Studies*.
- Sockin, J. (2024). Show me the amenity: Are higher-paying firms better all around.
- Sokolova, A. and T. Sorensen (2021). Monopsony in labor markets: A meta-analysis. *ILR Review* 74(1), 27–55.
- Sorkin, I. (2018, 01). Ranking Firms Using Revealed Preference. *The Quarterly Journal of Economics* 133(3), 1331–1393.
- Stansbury, A. and L. H. Summers (2020). The declining worker power hypothesis: An explanation for the recent evolution of the american economy. *Brookings Papers on Economic Activity*, 1–77.
- Stole, L. A. and J. Zwiebel (1996). Intra-firm bargaining under non-binding contracts. *The Review of Economic Studies* 63(3), 375–410.
- Syverson, C. (2019, August). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives* 33(3), 23–43.
- Van Reenen, J. (1996, 02). The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies. *The Quarterly Journal of Economics* 111(1), 195–226.
- Vilhuber, L. (2018, May). Lehd infrastructure s2014 files in the fsrdc. Working Paper CES WP-18-27, U.S. Census Bureau, Center for Economic Studies. Revision: May 8, 2018.
- Webber, D. A. (2022). Labor market competition and employment adjustment over the business cycle. *Journal of Human Resources* 57(S), S87–S110.
- Yeh, C., C. Macaluso, and B. Hershbein (2022, July). Monopsony in the us labor market. *American Economic Review* 112(7), 2099–2138.
- İmrohoroglu, A. and S. Tüzel (2014). Firm-level productivity, risk, and return. *Management Science* 60(8), 2073–2090.


Figures

Figure 1: Firm-Level Employment, Wage, and Productivity Growth Responses to a Firm-Specific Stock Return Shock




Note: This figures shows the growth responses of firm employment, wages, total factor productivity (from [İmrohoroglu and Tüzel \(2014\)](#)), and labor productivity (value-added per worker) to a stock return shock, as in (10) in the main text for $h = -5$ to 5 years. Confidence intervals are based off standard errors double clustered by industry and year.

Figure 2: Employment and Wage Responses to Stock Return Shock, by Worker Skill




Note: This figures shows the employment and wage responses from estimating (10) in the main text for $h = -5$ to 5 years for workers of different skill levels. Individuals in the bottom two quintiles of the cross-sectional distribution of worker effects are considered low-skilled, the third and fourth quintiles middle-skilled, and the top quintile high-skilled. Confidence intervals are based off standard errors double clustered by industry and year.

Figure 3: Estimated Supply Elasticities Trend Downward Over Time For All Skill Groups




Note: This figure shows estimates of supply elasticities implied by estimating employment and wage responses from (9) over time. The elasticity for a given year is obtained using a three-year moving window centered at that year (except for 1992 and 2013, which respectively use the 1991-1992 and 2012-2013 windows due to data availability). The sample spans 1992-2013. High skill workers are in the top quintile of individual fixed effects from a wage of individual earnings into worker- and firm-specific components; low skill workers are in the bottom two quintiles and middle skill the third and fourth quintiles.

Figure 4: Decline in Labor Supply Elasticities: Labor Productivity Sorts, 1991-2002 Versus 2003-2014



Note: This table reports estimates of estimates of the labor supply elasticity implied by estimating the employment and wage responses to stock returns, as in equation (13) of the main text, with estimates computed separately for the 1991-2002 and 2003-2014 subperiods. The main independent variable is the firm's own stock return in excess of the risk-free rate, and controls include 3-digit NAICS industry by year lagged growth rates in wages, employment, and total assets; the contemporaneous change in average worker skill level at the firm (see (A.8) for definition); and productivity quartile fixed effects. 95 percent confidence intervals are based off standard errors clustered by industry and year. See main text for further details.

Figure 5: The Cross-Sectional Spreads in Labor Shares and Estimated Markdowns Between High- and Low-Labor Productivity Firms Move Together Over Time



Note: This figure shows differences elasticity implied log markdowns between top- and bottom-quartile labor productivity firms from (15) in the main text, and log labor share differences following (16). Time-varying elasticity estimates for a given year are obtained using a three-year moving window centered at that year (except for 1992 and 2013, which respectively use the 1991-1992 and 2012-2013 windows due to data availability). Both series are standardized to unit standard deviation and zero mean. The sample period spans 1992-2013. The two series have a correlation of 0.73.

Tables

Table 1: Labor supply elasticity estimates: homogeneous and by worker skill type

	All	Low Skill	Middle Skill	High Skill
Employment	0.12 (0.010)	0.13 (0.011)	0.11 (0.011)	0.10 (0.010)
Wages	0.046 (0.004)	0.016 (0.002)	0.019 (0.001)	0.084 (0.005)
Elasticity	2.53 (0.30)	7.97 (0.88)	5.92 (0.59)	1.22 (0.17)

Note: This table shows estimates of the supply elasticities implied by the employment and wage responses to stock returns from estimating (9) in the text. Controls include 3-digit NAICS industry by year and productivity quartile fixed effects; lagged growth rates in wages, employment, and total assets; and the contemporaneous change in average worker skill level at the firm (see (A.8) for definition). Workers are grouped into skill groups based on their estimated worker effects from a modified Abowd et al. (1999) style wage decomposition with time-varying firm fixed effects. Individuals in the bottom two quintiles of the cross-sectional distribution of worker effects are considered low-skilled, the third and fourth quintiles middle-skilled, and the top quintile high-skilled. Changes in average worker skill are computed within the population of workers considered in the specification. Standard errors clustered by industry and year are in parentheses; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages.

Table 2: Productive firms have lower labor shares, higher valuations, and better operating performance, without high investment rates

Panel A: Labor Productivity and Firm Outcomes					
	log Lshare	Q (Tot)	log M/B	ROA _t	ROA _{t+1}
log VA/Worker	-0.56 (0.04)	0.64 (0.06)	0.25 (0.03)	0.11 (0.004)	0.08 (0.003)
Size Controls	X	X	X	X	X
Industry X Year FE	X	X	X	X	X
N	57500	53500	55000	55500	48000
R ² (within)	0.48	0.09	0.07	0.32	0.18

Panel B: Investment Rates and Tobin's Q by Labor Productivity Quartile

Productivity:	Quartile 1	Quartile 2	Quartile 3	Quartile 4	P-val 4-1
Inv Rate (Total)	0.18	0.16	0.16	0.20	0.33
Q (Total)	0.73	0.79	1.04	1.70	0.00

Note: Panel A of this table shows the relationship between log value added per worker and other firm-level outcomes. The variable log Lshare is the firm labor share measure described in the main text; log Q (Tot) is the log of intangible-adjusted total Q from [Peters and Taylor \(2017\)](#); log M/B is the log market/book ratio; and, ROA is the return on assets, defined as operating income before depreciation over total capital (the sum of physical capital and [Peters and Taylor \(2017\)](#) total intangible capital). Operating performance and valuation ratios are winsorized at the 1% level by year. See section 3.1 in main text for more details. Standard errors double clustered by year and industry in parentheses. Panel B of this table shows average investment rates in capital and in new hires for firms of different productivity quartiles. "P-val 4-1" gives the p-value from a test that the coefficients for firms in the top and bottom quartiles have equal values.

Table 3: Productive firms face lower supply elasticities for workers of all skill levels

Productivity:	Quartile 1	Quartile 2	Quartile 3	Quartile 4	P-val 4-1	R-sq	N
Panel A: Whole Firm							
Employment	0.12 (0.01)	0.09 (0.01)	0.12 (0.01)	0.11 (0.01)	0.249	0.13	43500
Wages	0.03 (0.003)	0.03 (0.003)	0.05 (0.004)	0.09 (0.01)			
Elasticity	4.32 (0.87)	2.71 (0.43)	2.34 (0.28)	1.17 (0.15)	0.001		
Panel B: Low Skill Workers							
Employment	0.13 (0.01)	0.10 (0.01)	0.14 (0.01)	0.14 (0.02)	0.522	0.11	43500
Wages	0.01 (0.002)	0.02 (0.003)	0.02 (0.003)	0.03 (0.003)			
Elasticity	10.51 (3.04)	6.37 (0.90)	6.77 (1.08)	4.66 (0.39)	0.042		
Panel C: Middle Skill Workers							
Employment	0.12 (0.02)	0.08 (0.01)	0.12 (0.01)	0.10 (0.01)	0.307	0.09	43500
Wages	0.01 (0.002)	0.02 (0.002)	0.02 (0.002)	0.03 (0.003)			
Elasticity	8.20 (1.70)	5.30 (0.79)	5.30 (0.71)	3.55 (0.44)	0.014		
Panel D: High Skill Workers							
Employment	0.10 (0.01)	0.08 (0.01)	0.11 (0.01)	0.09 (0.01)	0.148	0.08	43500
Wages	0.05 (0.01)	0.06 (0.004)	0.08 (0.01)	0.12 (0.01)			
Elasticity	2.04 (0.38)	1.39 (0.23)	1.39 (0.17)	0.74 (0.11)	0.004		

Note: This table contains supply elasticity estimates for firms sorted on log value-added/worker quartiles as in (13) in the text. Controls include 3-digit NAICS industry by year and productivity quartile fixed effects; lagged growth rates in wages, employment, and total assets; and the contemporaneous change in average worker skill level at the firm (see (A.8) for definition). Workers are grouped into skill groups based on their estimated worker effects from a modified Abowd et al. (1999) style wage decomposition with time-varying firm fixed effects. Individuals in the bottom two quintiles of the cross-sectional distribution of worker effects are considered low-skilled, the third and fourth quintiles middle-skilled, and the top quintile high-skilled. Changes in average worker skill are computed within the population of workers considered in the specification. "P-val 4-1" gives the p-value from a test that the coefficients for firms in the top and bottom quartiles have equal values. Wage data are from the LEHD, and the sample period spans 1991-2014. Standard errors clustered by industry and year are in parentheses; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages. See section 3 in main text for more details.

Table 4: Productive firms face lower supply elasticities: alternative shocks to labor demand

Productivity:	Below Median	Above Median	P-val (Above - Below)	P-val (J-Test)
Elasticity Estimates—Customer Sample				
Customer Ret	3.40 (1.47)	1.45 (0.52)	0.106	
Baseline	2.86 (0.55)	1.47 (0.24)	0.008	
Elasticity Estimates—Main Sample				
Patents	2.99 (0.65)	1.97 (0.69)	0.038	
Earnings Ret	3.34 (0.66)	1.87 (0.29)	0.011	
Baseline	3.53 (0.57)	1.67 (0.17)	0.000	
Elasticity Estimates—Overidentified Specifications				
Customer Ret	2.88 (0.54)	1.46 (0.23)	0.006	0.86
Patents	3.45 (0.58)	1.69 (0.17)	0.001	0.62
Earnings Ret	3.49 (0.59)	1.64 (0.16)	0.001	0.45

Note: This table shows supply elasticity estimates using alternative labor demand shifters. Firms are sorted into above- and below-median labor productivity bins. See appendix section [IA.6.3](#) for further details on these variables. All specifications have the baseline controls from (9) in main text, including industry $\times$ year fixed effects. “Baseline” refers to my main specification using one-year excess stock returns as a labor demand shock; for comparison in like samples, I estimate baseline elasticities for the set of firms where I can observe customers’ stock returns (panel A), and for the main sample (panel B). “P-Value (Above - Below)” gives the p-value on the differences in elasticities between the above- and below-median productivity firms. Standard errors double clustered by industry and year are in parentheses, and are computed by estimating the elasticity via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages.

Table 5: Elasticity-implied firm-level log wage markdowns are associated with firm labor shares, valuations, and profits

	Log Lshare		Q (Tot)		ROA_t		ROA_{t+1}	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Markdown	0.90 (0.11)	0.69 (0.10)	-1.50 (0.15)	-0.77 (0.12)	-0.26 (0.022)	-0.17 (0.019)	-0.20 (0.019)	-0.083 (0.011)
Log Markup	-0.28 (0.078)	-0.48 (0.135)	0.95 (0.16)	0.67 (0.16)	0.071 (0.018)	0.11 (0.031)	0.056 (0.014)	0.055 (0.015)
Size Controls	X	X	X	X	X	X	X	X
Capital Intensity Controls	X	X	X	X	X	X	X	X
Ind $\times$ Year FE	X	X	X	X	X	X	X	X
Firm FE		X		X		X		X
N	55500	55000	52000	51500	55500	55000	52000	51000
R^2 (within)	0.30	0.28	0.18	0.16	0.34	0.36	0.22	0.13

Note: This table reports results from regressing firm log labor shares, valuations, and contemporaneous and future return on assets on elasticity-implied log wage markdowns and log markups. Implied firm wage markdowns are computed by taking elasticity estimates from the top panel of Table 3. Capital intensity controls include the logs of physical and intangible capital per worker, while size controls include the logs of assets, sales, and employment. Markups are computed as in [De Loecker et al. \(2020\)](#). Standard errors clustered by industry and year in parentheses.

Table 6: Valuing labor market power: the value of wage markdowns in the aggregate and per worker

	Pooled firm-level average:				Labor market profit as share of aggregate:		
	Profits per Worker (\$)	Net inverse Markdown	Implied Elasticity	Empirical Elasticity	Capital Income	Capital Income (Adj.)	Revenues
Panel A: Homogeneous							
Overall Average	17,550	1.22	4.48	2.53	0.25	0.22	0.039
Overall Average, 1991–2002	10,570	1.16	6.17	2.92	0.20	0.17	0.029
Overall Average, 2003–2014	25,840	1.30	3.39	1.88	0.31	0.28	0.050
Low Productivity	4,891	1.09	10.63	3.53	0.13	0.11	0.011
High Productivity	30,210	1.35	2.84	1.67	0.28	0.25	0.052
Panel B: Skill Heterogeneity							
Overall Average	17,310	1.20	4.90	2.53	0.23	0.20	0.036
Low Skill	3,923	1.14	6.94	7.97	0.017	0.015	0.0026
Middle Skill	8,678	1.16	6.10	5.92	0.059	0.052	0.0091
High Skill	39,910	1.23	4.44	1.22	0.16	0.14	0.025

Note: This table uses firm-level estimated net inverse markdowns Π_t (defined in (19) and following the procedure described in internet appendix section IA.1.4). Profits per worker are reported in real 2011 US dollars, and are given by $(\Pi_t - 1) \times W_t$, where W_t is LEHD spending per worker. Implied elasticities are calculated as $1/(\bar{\Pi} - 1)$, where $\bar{\Pi}$ is the average of Π_t for the given set of firms. Panel A imposes the same Π_t for all workers within a firm, while Panel B allows for a different Π_t by worker skill, as defined in Table 1. The last three columns sum total profits from labor market power across all firms in a given year, and compare them to the total capital income (operating income before depreciation in that year); adjusted capital income (capital income plus expensed non-labor intangible investments); and total revenues for the given set of firms or time period. See section 5 of the main text and appendix section IA.1.4 for more details.

Table 7: Differences in labor shares explained by markdowns: cross-section and time series

Panel A: Cross-Sectional Labor Share Gap				
	Productivity		High – Low	% Explained
	Low	High		
Average Log Labor Share	-0.30	-0.73	-0.43	
Log Markdown (Homogeneous)	-0.09	-0.30	-0.21	50%
Log Markdown (Skill Heterogeneity)	-0.11	-0.26	-0.15	35%
Panel B: Time Series Change in Labor Share (1991 to 2014)				
	Year		Change	% Explained
	1991-2002	2003-2014		
Aggregate Labor Share	0.579	0.524	0.055	
<i>Counterfactual Aggregate Labor Share:</i>				
Homogeneous		0.544	0.036	35%
Skill Heterogeneity		0.558	0.021	61%

Note: Panel A of this table decomposes the fraction cross-sectional average labor share differences between high- and low-productivity firms that can be attributed to wage markdowns. In panel A I use markdown estimates for the full 1991-2014 period reported in Table 6. Panel B of this table decomposes the time series change in the aggregate labor share between the 1991-2002 and 2003-2014 period using markdowns estimated separately for the two subperiods. The counterfactual aggregate labor share is calculated by computing the counterfactual labor share if wage markdowns were generated from supply elasticities and adjustment costs parameters held constant at their 1991-2002 values in the 2003-2014 period. See section 5 of the main text for further details.

Table 8: Fraction of cross-sectional spreads in return on assets and valuations explained by wage markdowns

		Counterfactual Type:	
	Actual	Homogeneous	Skill Heterogeneity
Panel A: Return on Assets			
Low productivity	0.095	0.070	0.066
High productivity	0.165	0.099	0.110
High - Low	0.070	0.029	0.038
% Explained		58%	36%
Panel B: Firm Valuations (Total Q)			
Low productivity	0.77	0.67	0.64
High productivity	1.38	0.95	1.02
High - Low	0.61	0.28	0.38
% Explained		54%	38%

Note: Panel A of this table examines the counterfactual firm return on assets (defined as operating income before depreciation over physical plus intangible capital) by firm productivity rank if the dollar value of wage markdowns were paid to workers instead of the firm. Panel B of this table estimates predicted counterfactual valuations when the dollar value of wage markdowns is paid to workers instead of the firm. Counterfactual valuations are obtained from cross-sectional regressions of the [Peters and Taylor \(2017\)](#) total Q valuation on return on assets and investment rates and then using regression coefficients to predict the valuation that would obtain when markdown profits are paid to workers. See appendix section [IA.2](#) for further details on this procedure.

7 Main Appendices

A Data Appendix

A.1 Constructing Labor Productivity, Wages, Labor Shares, and Other Firm-Level Variables

Firm Value-Added and Labor Shares

I follow [Donangelo et al. \(2019\)](#) in defining value-added for Compustat firms following as the sum of operating income before depreciation, changes in inventories, and labor expenses.

$$VA_{j,t} = OIBDP_{j,t} + \Delta INVFG_{j,t} + LABEX_{j,t} \quad (A.1)$$

Changes in inventories are set to zero when missing. Here

$$LABEX_{j,t} = \widetilde{W}_{j,t} \times (EMP_{j,t} + EMP_{j,t-1})/2 \quad (A.2)$$

Instead of imputing wages as industry-size cell averages of Compustat item *XLR* as in [Donangelo et al. \(2019\)](#), I create a measure of the average wage ($\widetilde{W}_{j,t}$) paid to workers using my LEHD-Compustat match. I detail my computation of the average wage later on in this section. Because coverage of the LEHD varies by year and may not contain all establishments from a given Compustat gvkey in a given year, I multiply the inferred LEHD wage by the average of the firm's employment in years t and $t - 1$, rather than directly summing up LEHD wage compensation. Because Compustat employment is reported at year-end, I also follow [Donangelo et al. \(2019\)](#) in taking the average employment in adjoining years. Labor productivity is given by

$$\log(VA/Worker)_{j,t} = \log(VA_{j,t}/EMP_{j,t}) \quad (A.3)$$

Unless otherwise specified, I refer hereafter to labor productivity, productivity, and log value-added per worker interchangeably. Finally, I define the labor share of firm i at time t by

$$LSHARE_{j,t} = \frac{LABEX_{j,t}}{VA_{j,t}} \quad (A.4)$$

AKM Wage Decomposition

I now detail how I decompose wages into worker- and firm-specific heterogeneity in the tradition of [Abowd et al. \(1999\)](#) (AKM). I start with a modification of the AKM decomposition proposed by [Lachowska et al. \(2020\)](#) and [Engbom and Moser \(2020\)](#) that allows for the firm-specific component of wages to vary by time. Let i index individual workers; $j(i, t)$ a function indicating the firm j that employs individual i at time t ; and $X_{i,t}$ a third-degree

polynomial in worker age that is flat at age 40, as in [Sorkin \(2018\)](#) and [Card, Heining, and Kline \(2013\)](#). I then estimate the wage decomposition

$$\log(w)_{ijt} = \alpha_i + \phi_{j(i,t),t} + \beta X_{i,t} + \epsilon_{i,t} \quad (\text{A.5})$$

I estimate (A.5) for matched Compustat firms for overlapping 5-year moving windows. Because my analysis is at the annual frequency and LEHD earnings are quarterly, each wage w_{ijt} represents a full-year equivalent real wage for individual i in year t , as in [Sorkin \(2018\)](#). In order to be included in the sample for estimating (A.5), firm j must be worker i 's primary employer for that year (the firm with highest earnings), worker i must have been employed at that firm for at least two consecutive quarters within the year, and the worker must have earned more than \$3250 in 2011 US dollars. Papers performing AKM wage decompositions on LEHD earnings data set a lower threshold on annual earnings because earnings because hours worked are not observed in the LEHD. Due to the large size of the LEHD data, I estimate (A.5) for four disjoint 25% subsamples of my original sample of individuals found in the LEHD. All firm level aggregates taken from these estimates represent averages across these four 25% subsamples. Estimating (A.5) also requires that firm-years and individual worker pairs must belong to a set satisfying connectedness conditions in order for the fixed effects α_i and $\phi_{j(i,t),t}$ to be separately identified. I provide more details on this connectedness requirement and the sampling procedure in appendix A.

I use the sample of workers in (A.5) to create my inferred firm average wage $\widetilde{W}_{j,t}$. Let $N_{j,t}^{ins}$ denote the total number of workers mapped to firm j that are in the sample of (A.5) in year t . Let $N_{j,t}^{not-ins}$ denote the number of unique individuals that show up on the payrolls of firm j in year t that are not included in the sample (either due to earnings below the required threshold or firm j not being the primary employer for that year). Let $\widetilde{w}_{ijt}$ denote the actual (not full-year adjusted) year t earnings of worker i at firm j . Then I compute the firm-specific wage as

$$\widetilde{W}_{j,t} = \frac{\sum_i \widetilde{w}_{ijt}}{N_{j,t}^{ins} + N_{j,t}^{not-ins}} \quad (\text{A.6})$$

I contrast this unadjusted wage $\widetilde{W}_{j,t}$ with a full-time full year equivalent adjusted wage $W_{j,t}$. Let $\Gamma_{j,t}$ denote the set of workers that are in the sample of (A.5) for firm j in year t :

$$W_{j,t} = \frac{\sum_{i \in \Gamma_{j,t}} w_{ijt}}{N_{j,t}^{ins}} \quad (\text{A.7})$$

I use the actual average earnings $\widetilde{W}_{j,t}$ multiplied by Compustat employment for computing firm level value-added, labor expenses, and labor shares, as in (A.1), (A.2), and (A.4). Meanwhile, I focus on the employment response to log changes in the firm's full-time, full-year equivalent

adjusted wage $W_{j,t}$ when I estimate supply elasticities in Section 3. I make these choices in order to ensure that my computation of labor expenses represents actual spending on labor per Compustat employee, while estimated supply elasticities correspond to the employment response induced by a change in the wage offer for a consistently defined period of employment.

Firm-Level Skill Measure Derived From AKM Estimates

I introduce a measure of firm average worker skill that I derive from the wage decomposition in (A.5). This is given by:

$$s_{j,t} = \log \left(\frac{\sum_{i \in \Gamma_{j,t}} \exp(\hat{\alpha}_{i,t})}{N_{j,t}^{ins}} \right) \quad (\text{A.8})$$

Thus $s_{j,t}$ is the log of the average component of worker wages coming from the worker-specific heterogeneity of the firm's labor force. Hence firms with more skilled workers will have a higher $s_{j,t}$. When I estimate specifications for incumbents, recruits, or specific skill levels, I recalculate $s_{j,t}$ by restricting to only workers within the firm from the given category.

A.2 Sampling/Cleaning LEHD Wage Data and Estimation of Wage Decomposition

The basic person-level identifier variable in the LEHD data is the PIK, which has a one-to-one correspondence with an individual's Social Security number. My baseline LEHD sample comes from the list of unique PIK identifiers obtained from the union of all individuals found in the 2000 Decennial Census; the SIPP and Current Population Survey; and a 10% subsample of the Numident sample. I further generate four disjoint 25% subsamples of this list of PIKs intersected with the list of PIKs in the LEHD.

I repeat the following steps for each of the four disjoint 25% subsamples. All estimates are taken separately across these disjoint subsamples; all firm level aggregates represent averages across these four subsamples. For each year in the LEHD 1990-2015 I retain all individuals who were found to be employed at a Compustat-linked firm in that year. This forms my LEHD-Compustat match. I then clean LEHD earnings data using a procedure that follows very closely with Sorkin (2018). Because days and hours worked in the LEHD are not observed, these steps are meant to convert quarterly LEHD earnings to their full-year wage equivalents. I categorize quarterly earnings observations into three groups: full, continuous, and discontinuous. Quarter q is a full earnings occur when individual i is linked to firm j at quarters q , $q - 1$ and $q + 1$; quarter q is a continuous is when individual i is linked to firm j at quarters q and on of $q - 1$ and $q + 1$, but not both; finally, discontinuous quarters occur when individual i is linked to firm j at during quarter q but not at $q - 1$ or $q + 1$.

Because full quarters are most likely to represent full-time employment, they are prioritized as follows. If individual i has any full quarters of employment at firm i in the given year,

then annual wages are taken to be

$$4 \times \text{Total Earnings in Full Quarters} / \text{Number of Full Quarters}.$$

If there are no full quarters of employment, then annual earnings are

$$8 \times \text{Total Earnings in Continuous Quarters} / \text{Number of Continuous Quarters}.$$

Finally, if there are no continuous quarters, annual earnings are given by

$$12 \times \text{Total Earnings in Discontinuous Quarters} / \text{Number of Discontinuous Quarters}.$$

Assuming separations occur uniformly within a quarter, a continuous quarter represents a half quarter of employment at the firm and a discontinuous quarter represents a third of a quarter of employments, so this adjusts wages to full-year terms. Full quarters require no such adjustment, which is the reason for prioritizing full quarters of employment over others. Earnings are further adjusted to real equivalents. I follow [Sorkin \(2018\)](#) in using the CPI from the 4th quarter of 2011 as the baseline for this real wage adjustment.

I link each individual to their primary in each year. The primary employer is the one where their unadjusted earnings are highest. I define the employer at the gvkey level for worker-firm-years linked to a Compustat firm, and at the EIN level for worker-firm-years with no such link. Since I restrict to persons linked to Compustat in the current year, every individual in my sample will have at least one job in year t linked to a Compustat gvkey, although this will not always be their primary employer. Using the adjusted earnings I estimate the modified AKM decomposition for individuals' primary employers:

$$\log(w)_{ijt} = \alpha_i + \phi_{j(i,t),t} + \beta X_{i,t} + \epsilon_{i,t} \quad (\text{A.9})$$

As in [Sorkin \(2018\)](#), I require individuals to have earned more than \$3250 in real earnings at their employer in the year in order to be considered part of the sample for estimating (A.9). Since hours or days worked are not observed, this drops individuals likely to have had a minimal attachment to the firm in the year due to part-time employment. I estimate (A.9) for overlapping 5-year periods starting in 1991 and ending in 2010. Because I have estimates for overlapping 5-year intervals, I take estimates for the final year from the sample (i.e. firm and person effects in 2013 come from the 2009-2013 subsample, 2014 comes from 2010-2014). The exception is for 1991-1994, which do not have a full 5-year period ending in the given year, so estimates for these years come from the 1991-1995 subsample. This gives me estimates of model parameters for the 1991-2014 period (though the LEHD covers 1990 and 2015, I use these years to determine full/continuous/discontinuous quarters in 1991 and

2014, respectively, leading my sample period to span 1991 to 2014).

Fixed effects in (A.9) are only defined for the collection of firm-years connected by worker flows across firms and time. Because I have time-varying firm effects, the connectivity requirements are slightly different than in the classic AKM decomposition with time-invariant firm effects (see Lachowska et al. (2020) and Engbom and Moser (2020) for a detailed discussion on the connectivity requirements in order for parameters to be identified). Accordingly, when estimating (A.9) I restrict the sample to the largest connected set of worker-firm-year observations following these two papers. Because Compustat firms are large, in practice this connectivity restriction drops a miniscule fraction of the data.

A.3 Imputing Compustat Staff and Labor Expense (XLR)

Here I describe my imputation of Compustat variable XLR . I use $LABEX_{j,t}$ from (A.2) in the main text. I use $LABEX_{j,t}$ to predict XLR out-of-sample (I also use predicted XLR even when actual XLR is available for consistency). To ensure the predicted XLR is positive I run the following regression in logs:

$$\log(XLR_{j,t}) = \alpha + \alpha_{I(j),t} + \beta_t \log(LABEX_{j,t}) + \epsilon_{j,t} \quad (\text{A.10})$$

Here $\alpha_{I(j),t}$ are two-digit NAICS by year fixed effects and β_t are time-varying coefficients on $\log(LABEX_{j,t})$.

My predicted level of XLR is then just $\widehat{XLR}_{j,t} = \exp(\hat{\alpha} + \hat{\alpha}_{I(j),t} + \hat{\beta}_t \log(LABEX_{j,t}))$. When estimation error is taken into consideration, simply taking the exponential of the predicted log of the variable in question often leads to better forecasts (Brardsen and Lütkepohl, 2011). A Jensen's inequality adjustment term for a scaling factor, using the variance of the residuals and assuming lognormality leads to a non-trivial overestimate of the actual XLR , while the average level of $\widehat{XLR}_{j,t}$ is much closer to actual $XLR_{j,t}$. Because of this, I use the exponential of the log to create a strictly positive predicted value of $XLR_{j,t}$. In-sample $\widehat{XLR}_{j,t}$ has a correlation of 0.95 with actual $XLR_{j,t}$. When I analyze trends aggregate labor share, I take aggregate labor expenses implied by summing up $\widehat{XLR}_{j,t}$ across firms, since it generates a public-firm labor share decline that is more in line with other estimates (Karabarbounis and Neiman, 2013; Hartman-Glaser et al., 2019). See also appendix Figure IA.4 for a comparison of aggregate labor share dynamics using implied aggregate labor shares implied by using $\widehat{XLR}_{j,t}$ versus $LABEX_{j,t}$ to sum up aggregate labor expenses. Cross-sectionally, using $LABEX_{j,t}$ or $\widehat{XLR}_{j,t}$ makes little difference, so I use the more direct $LABEX_{j,t}$ measure to construct labor productivity measures and when analyzing log labor share variation across firms.

8 Internet Appendix

IA.1 Dynamic Monopsony Framework

IA.1.1 Example dynamic model

Consider the following extension of the example static monopsony framework from section 2.2.1 in the main text. In period t the firm draws a mass of Γ_t workers who consider the firm's wage offer. The firm faces an inverse labor supply curve of the following form:

$$W(L_t; \Gamma_t) = \left(\frac{L_t}{\Gamma_t} \right)^{\frac{1}{\beta}} \times (\bar{W} - b) + b \quad (\text{IA.1})$$

This is the same functional form as in [Kline et al. \(2019\)](#) and [Card et al. \(2018\)](#), who both assume labor supply is proportional to $(W - b)^\beta$ for some parameters b and β . To see this, express in terms of wages:

$$L(W_t; \Gamma_t) = \Gamma_t \times \left(\frac{W_t - b}{\bar{W} - b} \right)^\beta \quad (\text{IA.2})$$

Then $\left(\frac{W_t - b}{\bar{W} - b} \right)^\beta$ is the probability that a given worker accepts employment at the firm at the wage offer W_t . As [Card et al. \(2018\)](#) note, the parameter b acts like an outside option, and β captures the ease of substitution across firms or worker bargaining power.

The supply elasticity under this functional form is

$$\frac{\partial L_t}{\partial W_t} \frac{W_t}{L_t} = \frac{\beta W_t}{W_t - b} \quad (\text{IA.3})$$

This is decreasing in the wage offer. As more productive firms endogenously choose a higher wage, this ensures supply elasticities decrease with productivity, a key regularity in the data.

I assume that the process for Γ_t follows $\Gamma_t = \mathcal{L}_{t-1} + L_{t-1}$, where L_{t-1} is the number of workers at the firm at the end of the last period, and $\mathcal{L}_{t-1}$ is the mass of workers from the external market who draw the firm. Because the number of workers the firm draws Γ_t is increasing in the number of workers at the firm last period, the firm is constrained in how fast it can grow, and so it accumulates workers slowly over time. Although firm hiring increases Γ_t , I assume that the firm views itself as “intertemporally small”, in the sense that it takes Γ_t as given each period, and hence does not internalize the dynamic impact of its hiring decisions on the size of the pool of workers it will draw in the future.²⁹

²⁹This assumption simplifies the firm's Euler equation considerably, and allows me to avoid having to estimate an additional “intertemporal labor supply elasticity” parameter when I eventually take the framework to the data. What's more, I find that if the firm optimizes over employment and fully internalized the impact on Γ_t in the Euler equation, it would generate markdowns that are negative, correlate strongly negatively with labor productivity, and decline over time—all highly inconsistent with empirical patterns in profitability, valuations, and labor shares. And finally, it's not obvious what is the “correct” way one would incorporate

On the other hand, the firm fully internalizes the impact of its labor demand on future adjustment costs. I assume the firm uses both labor and capital in production, and both are costly to adjust. Labor at time t can be adjusted at cost

$$\Phi_L(H_t, L_{t-1}) = \frac{\gamma_L}{2} \left(\frac{H_t(L_t, L_{t-1})}{L_{t-1}} \right)^2 L_{t-1} \quad (\text{IA.4})$$

Here $H_t(L_t, L_{t-1})$ are time- t gross hires. Gross hires are in turn a direct function of past employment L_{t-1} , the time- t total labor demand L_t , and the exogenous outside pool $\mathcal{L}_{t-1}$:

$$H_t \equiv \lambda_{t-1}(L_{t-1}, \mathcal{L}_{t-1}) \times L_t, \quad (\text{IA.5})$$

where

$$\lambda_{t-1}(L_{t-1}, \mathcal{L}_{t-1}) \equiv \frac{\mathcal{L}_{t-1}}{\mathcal{L}_{t-1} + L_{t-1}} \quad (\text{IA.6})$$

is the share of new hires in total employment L_t . The direct dependence of H_t on L_t and L_{t-1} ensures that the firm has a single control which determines its optimal labor demand.

Firm capital depreciates at a rate δ , and capital can be adjusted at cost

$$\Phi_K(K_t, K_{t-1}) = \frac{\gamma_K}{2} \left(\frac{K_t - (1 - \delta)K_{t-1}}{K_{t-1}} \right)^2 K_{t-1} \quad (\text{IA.7})$$

Firms have a Cobb-Douglas production function in capital K_{t-1} and labor L_t , with idiosyncratic (log) productivity z_t :

$$F(L_t, K_{t-1}, z_t) = \exp(z_t + \mu_z) L_t^{1-\alpha} K_{t-1}^\alpha \quad (\text{IA.8})$$

The firm has a discount rate ρ , and define investment $I_t \equiv K_t - (1 - \delta)K_{t-1}$. The firm problem is then:

$$\max_{L_t, K_t} \sum_{t=0}^{\infty} \rho^t E_t [F(L_t, K_{t-1}, z_t) - W(L_t; \Gamma_{t-1})L_t - I_t - \Phi_K(K_t, K_{t-1}) - \Phi_L(H_t, L_{t-1})] \quad (\text{IA.9})$$

such dynamic interdependence into the firm's problem in the first place. When the supply curve is static it is immaterial whether the firm chooses labor subject to the inverse labor supply curve, or wages subject to the labor supply curve. However, if the firm internalizes that its labor demand today affects the level of the curve tomorrow, the optimal control when firms optimize over wages is quite different than when firms optimize over labor, respectively generating larger and smaller markdowns relative to static monopsony. For all these reasons it makes sense to model the firm as not internalizing this margin.

Taking first-order conditions, the Euler equation for L_t is:

$$\frac{\partial F(L_t, K_{t-1})}{\partial L_t} - \lambda_{t-1} \frac{\partial \Phi_L(H_t, L_{t-1})}{\partial H_t} = W(L_t; \Gamma_{t-1}) + \frac{\partial W_t(L_t; \Gamma_{t-1})}{\partial L_t} L_t + \rho E_t \left[\frac{\partial \Phi_L(H_{t+1}, L_t)}{\partial L_t} \right] \quad (\text{IA.10})$$

Where I have used that $\lambda_{t-1} \frac{\partial \Phi_L(L_t, L_{t-1})}{\partial H_t} = \frac{\partial \Phi_L(H_t, L_{t-1})}{\partial L_t}$, and

$$\frac{\partial W(L_t; \Gamma_{t-1})}{\partial L_t} = (1/\beta) \times (\bar{W} - b) \times L_t^{1/\beta-1} \times \Gamma_{t-1}^{-1/\beta} \quad (\text{IA.11})$$

$$\frac{\partial \Phi_L(H_t, L_{t-1})}{\partial H_t} = \gamma_L \left(\frac{H_t}{L_{t-1}} \right). \quad (\text{IA.12})$$

The final derivative inside the expectation on the right hand side of (IA.10) is

$$\frac{\partial \Phi_L(H_{t+1}, L_t)}{\partial L_t} = -\gamma_L \left(\frac{1}{2} + (1 - \lambda_t) \right) \times \left(\frac{H_{t+1}}{L_t} \right)^2, \quad (\text{IA.13})$$

where (IA.13) follows after applying the chain rule to account for the dependence of $H_{t+1} \equiv \frac{\mathcal{L}_\square}{\mathcal{L}_\square + L_t} \times L_{t+1}$ on L_t , and recognizing that $\frac{L_t}{\mathcal{L}_t + L_t} = (1 - \lambda_t)$.

The Euler equation for capital K_t is

$$\rho E_t \left[\frac{\partial F(L_{t+1}, K_t)}{\partial K_t} \right] - \frac{\partial \Phi_K(K_t, K_{t-1})}{\partial K_t} = 1 + \rho E_t \left[\frac{\partial \Phi_K(K_{t+1}, K_t)}{\partial K_t} \right] \quad (\text{IA.14})$$

where

$$\frac{\partial \Phi_K(K_t, K_{t-1})}{\partial K_t} = \gamma_K \left(\frac{I_t}{K_{t-1}} \right) \quad (\text{IA.15})$$

$$\frac{\partial \Phi_K(K_{t+1}, K_t)}{\partial K_t} = -\frac{\gamma_K}{2} \left(\frac{I_{t+1}}{K_t} \right)^2 - \gamma_K (1 - \delta) \left(\frac{I_{t+1}}{K_t} \right) \quad (\text{IA.16})$$

and the production function derivatives take the standard Cobb-Douglas form.

IA.1.2 Dynamic model calibration

I now explore the ability of this framework to speak to dynamic responses of firm quantities in the data, which requires a model calibration. I start by calibrating $\rho = \exp(-0.065)$, which maps to a real discount rate of approximately 6.5% (nominal discount rate of 9% minus 2.5% annual inflation in my sample period), and I take $\delta = 0.18$ to match the average investment rate of 18% for Peters and Taylor (2017) total capital that includes both physical and intangible capital. Next, I normalize $\mathcal{L}_t = 1$ and $\bar{W} = 1$. I assume the firm starts in a steady-state at the initial (log) TFP level μ_z , and I simulate a small (1%) deterministic mean-reverting shock z to TFP with persistence $\rho_z = 0.825$, which exactly matches the slope of the firm TFP response to stock returns in Figure 1. I calibrate the remaining parameters

μ_z , β , b , γ_L , γ_K , and α to roughly match the dynamic responses of employment, wages, and capital exhibited in Figure 1 and appendix Figure IA.2. I also target the initial supply elasticity of 2.5 in column (1) of Table 1; the average firm-level log labor share of -0.51; a worker separations rate of 30%; and, the ratio of the initial wage-to-TFP passthrough of 0.27.

The resulting calibration is $\beta = 0.32$; $b = 0.705$; $\gamma_L = 1.5$; $\gamma_K = 6.5$; $\alpha = 0.165$; and $\mu_z = 0.57$. The model delivers a labor elasticity of 2.5, along with a log labor share of -0.5; a separations rate of 0.29; and, a wage / TFP passthrough of 0.27, all extremely close to their empirical counterparts. Besides this accurate fit, the model captures the dynamics of employment, wages, and capital quite well, as can be seen in appendix Figure IA.3, where I plot the dynamics of the cumulative growth rates in average labor, wages, capital, and TFP in the model and in the data; for ease of comparison, I scale both model and data moments by dividing them by the size of the initial TFP response.

Because the firm gradually accumulates workers and is subject to both labor and capital adjustment costs, the model generates a slow increase in employment that is in line with the data, combined with an immediate spike in wages followed by a subsequent decline. Relative to the data, the model delivers an initial capital response that is slightly too high, with a long-run slope that is just a little too low; it seems likely that a more involved functional form for capital adjustments (such as dependence on both the level and the change in the investment rate or other non-convexities) could deliver more nuanced dynamics. Still, the capital response is qualitatively accurate and also close quantitatively, and my primary focus is on the dynamics of labor demand anyway. Overall, the model provides a solid fit.

How much do dynamics adjust model-implied wage markdowns? Following the definition in equation (19), the model delivers a net inverse wage markdown Π_t of about 1.22 (meaning wages are about 82% of net marginal revenues). Since in a static setting the supply elasticities ϵ relate to the net inverse wage markdown by $\Pi_t = \frac{\epsilon+1}{\epsilon}$, a static model with supply elasticity of roughly 4.6 would deliver the same markdown estimates. Thus adjustment costs temper the exercise of market power considerably, with firms marking down wages as if their supply elasticity was nearly twice as large as it actually is. Interestingly, this is almost exactly the same value I obtain when estimating the firm's Euler equation directly from the data in section 5 of the main text. I detail the exact estimation procedure that allows me to arrive at this empirical conclusion in appendix section IA.1.4 below.

What does the model say about supply elasticity differences across the productivity distribution? My baseline supply elasticity estimates in the middle panel of Table 4 include an elasticity of 1.67 for above-median labor productivity firms and 3.53 for below-median firms. Holding all other model parameters the same, I calibrate $\mu_z^{Low-prod} = 0.45$ and $\mu_z^{High-prod} = 0.80$ to roughly match the spread in average [İmrohoroglu and Tüzel \(2014\)](#) TFP between high- and low-labor productivity firms. This yields model-implied $\epsilon_{Low-prod} = 3.8$

and $\epsilon_{High-prod} = 1.5$. Therefore the model also generates a quantitatively accurate supply elasticity gap between high- and low-labor productivity firms under a realistic cross-sectional spread in TFP.

IA.1.3 Extended dynamic monopsony framework with incumbent and recruit Heterogeneity

I now show how an extension on the framework where the firm can make separate offers to new hires and retained incumbents provides an additional first-order condition that is informative about γ_L . Specifically, the setting is the exact same as in section IA.1.1, except now the firm faces the supply curves

$$W_I(L_t^I; \Gamma_t^I) = \left(\frac{L_t^I}{\Gamma_t^I} \right)^{\frac{1}{\beta}} \times (\bar{W} - b) + b \quad (\text{IA.17})$$

for workers who were employed at the firm last period, and

$$W_H(H_t; \Gamma_t^H) = \left(\frac{H_t}{\Gamma_t^H} \right)^{\frac{1}{\beta}} \times (\bar{W} - b) + b \quad (\text{IA.18})$$

for workers hired from outside the firm. Here $\Gamma_t^H = \mathcal{L}_{t-1}$ and $\Gamma_t^I = H_{t-1} + L_{t-1}^I \equiv L_{t-1}$. Adjustment costs are again quadratic in gross hiring:

$$\Phi_L(H_t, L_{t-1}) = \frac{\gamma_L}{2} \left(\frac{H_t}{L_{t-1}} \right)^2 L_{t-1} \quad (\text{IA.19})$$

where $L_{t-1} = H_{t-1} + L_{t-1}^I$.

Since total labor L_t rises one-for-one with H_t , the Euler equation for new hires H_t is

$$\frac{\partial F(L_t, K_{t-1})}{\partial L_t} - \frac{\partial \Phi_L(H_t, L_{t-1})}{\partial H_t} = W_H(H_t; \Gamma_{t-1}^H) + \frac{\partial W_H(H_t; \Gamma_{t-1}^H)}{\partial H_t} H_t + \rho E_t \left[\frac{\partial \Phi_L(H_{t+1}, L_t)}{\partial L_t} \right] \quad (\text{IA.20})$$

And the Euler equation for incumbent workers:

$$\frac{\partial F(L_t, K_{t-1})}{\partial L_t} = W_I(L_t^I; \Gamma_{t-1}^I) + \frac{\partial W_I(L_t^I; \Gamma_{t-1}^I)}{\partial L_t} L_t^I + \rho E_t \left[\frac{\partial \Phi_L(H_{t+1}, L_t)}{\partial L_t} \right] \quad (\text{IA.21})$$

First, recall that for a firm with wage function W , we can re-write $W + \frac{\partial W}{\partial L} L = \frac{\epsilon_{L,W} + 1}{\epsilon_{L,W}} W$, where $\epsilon_{L,W}$ is the elasticity of labor supply to the firm. Then, subtracting (IA.20) from (IA.21), we get

$$\frac{\epsilon_I + 1}{\epsilon_I} W_t^I - \frac{\epsilon_H + 1}{\epsilon_I} W_t^H = \frac{\partial \Phi_L(H_t, L_{t-1})}{\partial H_t} = \gamma_L \left(\frac{H_t}{L_{t-1}} \right) \quad (\text{IA.22})$$

where ϵ_I and ϵ_H respectively give the firm-specific labor supply elasticities for retained incumbents L_t^I and new hires H_t . Because W_t^I , W_t^H , H_t , and L_{t-1} can be observed in the data, and the elasticities can be estimated, (IA.22) provides a useful moment condition to estimate γ_L .

IA.1.4 Details on estimation of dynamic monopsony framework

In section 4 I show that the firm's Euler equation can be written as

$$\Pi_t = \frac{1 + \epsilon_{L_t, W_t}}{\epsilon_{L_t, W_t}} + \frac{1}{W_t} \rho E_t \left[\frac{\partial \Phi_{t+1}}{\partial L_t} \right], \quad (\text{IA.23})$$

I now describe how I estimate the different ingredients of the firm's Euler equation (IA.23) in the data. As mentioned in the main text, I calibrate the real discount rate $\rho = \exp(-.065)$, corresponding to a nominal discount rate of 9% minus 2.5% inflation. The main text discusses the estimation of supply elasticities ϵ_{L_t, W_t} in detail. With quadratic adjustment costs I only require a calibration of the parameter γ_L , which maps hiring rates into marginal costs. With γ_L in hand, I can then back out $\frac{\partial \Phi_{t+1}}{\partial L_t}$, the last component of (IA.23).

Estimating adjustment cost parameter γ_L

As equation (IA.22) makes clear, to back out γ_L I need to know the firm's hiring rate H_t/L_{t-1} ; the wage paid to incumbent workers W_t^I ; the wage paid to new hires W_t^H ; the supply elasticity of incumbent workers ϵ_I ; and the supply elasticity of new hires ϵ_H .

Since I can't directly observe the Compustat employment by worker incumbent or recruit type, I follow the same logic as in section 2.2.3 of the main text when I analyzed supply elasticities by skill type. Specifically, I assume the number of workers of a type $k = I, H$ are proportional to their observed shares of workers observed in the LEHD. So if I observe that the firm j has $N_{j,t}$ workers in the LEHD in year t , and $N_{j,t}^I$ of those are incumbent workers (defined as those who were also employed at the firm last year), then I assume Compustat employment among incumbent workers is equal to $EMP_{j,t}^I = EMP_{j,t} \times N_{j,t}^I/N_{j,t}$. Similarly, employment of new recruits H is given by $EMP_{j,t}^H = EMP_{j,t} \times N_{j,t}^H/N_{j,t}$. I take $W_{j,t}^I$ to be the average full-time, full-year equivalent adjusted wage for the firm's incumbent workers, and do the same for $W_{j,t}^H$. I then estimate for $k = I, H$

$$\log EMP_{j,t+1}^k - \log EMP_{j,t}^k = \alpha_{q(j,t)} + \alpha_{I(j),t} \sum_{q=1}^2 \varepsilon_q^k \times \mathbf{1}(q(j,t) = q) \times (\log W_{j,t+1}^k - \log W_{j,t}^k) + \Psi X_{j,t} + \epsilon_{j,t+1}^k \quad (\text{IA.24})$$

where $q(j,t)$ denotes an indicator for the firm's labor productivity ranking (here, an indicator for being in the below median or above median group in that year). I instrument for $\mathbf{1}(q(j,t) = q) \times (\log W_{j,t+1}^k - \log W_{j,t}^k)$ using firm stock returns interacted with productivity rank dummies, as in the main text. The controls $X_{j,t}$ are the same as described in main text

equation (13), where I am now careful to restrict the firm skill composition control (based off the change in average AKM worker intercepts) only to the set of workers in the relevant group $k = I, H$.³⁰ I estimate (IA.24) separately by time period (1991-2002 and 2003-2014) and by incumbent and recruit status, resulting in the supply elasticity estimates $\varepsilon_{q,\tau}^k$ for $k = I, H$ and $\tau = 1991-2002, 2003-2014$.

With these elasticity estimates in hand, I then take (IA.22) to the data by estimating

$$\widetilde{W_{j,t+1}^I} - \widetilde{W_{j,t+1}^H} = \sum_{q=1}^2 \gamma_{L,q} \times \mathbf{1}(q(j,t) = q) \times \left(\frac{EMP_{j,t+1}^H}{EMP_{j,t}} \right) + \Psi X_{j,t} + \epsilon_{j,t+1} \quad (\text{IA.25})$$

Following the identification strategy used throughout the paper, I instrument for the empirical gross hiring rate $\left(\frac{EMP_{j,t+1}^H}{EMP_{j,t}} \right)$ using stock returns. The hiring rate is also winsorized cross-sectionally at the 1% level to minimize the impact of a few large outliers. The notation $\widetilde{W_{j,t+1}^k}$ denotes wages for type k with two adjustments. First, following (IA.22) I adjust for the supply elasticities by multiplying the wages $W_{j,t+1}^k$ by the model-implied elasticity adjustment factor $(\varepsilon_{q,\tau}^k + 1)/\varepsilon_{q,\tau}^k$. Additionally, skill composition differences between a firm's incumbents and recruits can affect the wage premium between incumbents and recruits for reasons unrelated to adjustment costs. I account for such compositional effects by further multiplying $W_{j,t+1}^k$ by an adjustment factor that converts $W_{j,t+1}^k$ to the implied wage if group k had the same skill level as the typical firms' new recruits in that same year.³¹ Additionally, since both incumbent and worker wages are in the regression, I now include separate controls for the changes in firm average AKM worker effects for both the firm's incumbents and recruits in $X_{j,t}$.

Applying estimated adjustment costs to firm Euler equation

With γ_L estimates in hand, I can now calculate $\frac{\partial \Phi_{t+1}}{\partial L_t}$. I apply the same firm setup as in my example dynamic model. Following (IA.13), and noting that $(1 - \lambda_t)$ in the model maps to the firm's incumbent share of workers at time $t + 1$ (which is pre-determined at time t in the model), I compute the data analogue of (IA.13):

$$\frac{\partial \Phi_L(EMP_{j,t+1}^H, EMP_{j,t})}{\partial EMP_{j,t}} = -\widehat{\gamma_{L,q(j,t)}} \times \left(\frac{1}{2} + \frac{EMP_{j,t+1}^I}{EMP_{j,t+1}} \right) \left(\frac{EMP_{j,t+1}^H}{EMP_{j,t+1}} \right)^2 \quad (\text{IA.26})$$

Note that this derivative is negative, meaning that firms dial back market power because they understand that hiring more workers today can help them avoid paying future adjustment

³⁰I also allow coefficients on controls to vary by productivity rank in the new recruits specification, which alleviates some numerical instability in the 2SLS procedure in the first half of the sample.

³¹Specifically, I take the cross-sectional average of firm's worker AKM intercepts for their recruits using a version of the measure defined in equation (A.8) that only considers new hires. Call this cross-sectional average $\bar{s}_{t+1}^H$. Denote $s_{j,t+1}^k$ as the firm level average AKM intercept for firm j 's workers of type $k = I, H$. Then I skill composition-adjust wages by multiplying $W_{j,t+1}^k$ by $\exp(\bar{s}_{t+1}^H - s_{j,t+1}^k)$.

costs.

The final step is to compute the empirical expectation

$$\hat{E}_t \left[\frac{\partial \Phi_L (EMP_{j,t+1}^H, EMP_{j,t})}{\partial EMP_{j,t}} \right] \quad (\text{IA.27})$$

which I map to the data by taking the empirical average within above- and below-median labor productivity bins, interacted with time period (1991-2002 and 2003-2014). Expected future marginal adjustment costs are also divided by current firm wages in the Euler equation. Because I use wages that are adjusted for worker skill composition when I estimate γ_L , I make the analogous compositional adjustments to the firm-level average wage $W_{j,t}$.

Allowing for Skill Heterogeneity

The above discussion has assumed a single worker type (potentially up to incumbent and recruit status when needed). In the quantification exercise of section 5 in the main text I also allow for skill heterogeneity. The extension to skill heterogeneity is straightforward. With multiple worker skill types s the firm now has the following first-order condition for each s :

$$\Pi_t^s = \frac{1 + \epsilon_{L_t, W_t}^s}{\epsilon_{L_t, W_t}^s} + \frac{1}{W_t^s} \rho E_t \left[\frac{\partial \Phi_{t+1}^s}{\partial L_t^s} \right], \quad (\text{IA.28})$$

Then (IA.28) can be estimated separately for each s following the exact same procedure as before, except replacing all homogeneous firm choice variables with their type- s analogues: hiring rates are now relative to the stock of type s workers; supply elasticities are now estimated separately for different skill types (including for type- s incumbents and recruits, which are needed to back out γ_L^s); and so on. Using this approach, I also estimate (IA.28) for low-, middle-, and high-skill workers, following the skill breakdown used in Tables 1 and 3. This approach generates a different firm-level net inverse wage markdown Π_t^s for each skill type s .

IA.2 Effects of Labor Market Power on Valuation Ratios

As discussed in the main text, valuation ratios reflect the marginal incentive to invest in capital and also the of economic rents. Crouzet and Eberly (2021) derive an analytical expression for the different components of the Tobin's Q ratio when a firm faces convex adjustment costs. With one type of capital in production this expression is:

$$Q = \frac{1}{r-g} \frac{\Pi}{K} - \frac{1}{r-g} R + q \quad (\text{IA.29})$$

Here Q is the empirically observable Tobin's Q ratio; $\frac{1}{r-g}$ is a Gordon growth discount term; $\frac{\Pi}{K}$ is the operating profit to capital ratio; R is a user cost of capital term which depends on

depreciation rates and adjustment costs; and, q is the marginal value of an additional unit of capital, or the marginal incentive to invest.

While (IA.29) holds exactly under specific assumptions, more generally it highlights the direct link between firm operating profits and the component of valuations that are unrelated to the investment opportunities. Based off this relationship, I devise a simple empirical test to examine how rents from labor market power affect the average valuation ratios across the productivity distribution. Specifically, I first predict the Peters and Taylor (2017) total Tobin's Q ratio using the firm operating income-to-total capital ratio in the following cross-sectional regression:

$$Q_{j,t} = \beta_t \times \left(\frac{OI_{j,t}}{K_{j,t}} \right) + \delta_t \times \frac{I_{j,t+1}}{K_{j,t}} + \alpha_t + \epsilon_{j,t} \quad (\text{IA.30})$$

I estimate (IA.30) period-by-period and within each productivity group (to ensure that the conditional mean estimate is exactly equal to each productivity group's empirical average). I assume that user costs of capital and depreciation rates are constant across firms, which are absorbed in the α_t term. Abstracting away from financial constraints, the marginal incentive to invest q should be linear in the investment rate; accordingly, I control for firm investment rates $\frac{I_{j,t+1}}{K_{j,t}}$. I then take the estimated $\widehat{\beta}_t$ and predict counterfactual valuation ratios $\widetilde{Q}_{j,t}$, using the implied counterfactual operating income after subtracting off the dollar value of wage markdowns, as in section 5:

$$\widetilde{Q}_{j,t} = \widehat{\beta}_t \times \left(\frac{\widetilde{OI}_{j,t}}{K_{j,t}} \right) + \widehat{\delta}_t \times \frac{I_{j,t+1}}{K_{j,t}} + \widehat{\alpha}_t \quad (\text{IA.31})$$

where $\widetilde{OI}_{j,t}$ is firm j operating income before depreciation minus profits from labor market power. As described in the main text, profits from labor market power are defined as profits per worker type i , summed across all workers of a given type.³²

I then use (IA.31) to predict counterfactual average Peters and Taylor (2017) average total Tobin's Q valuation ratios by labor productivity rank, which I compare to their actual averages in Table 8.

IA.3 Elasticity Estimate Bias in a Simple Model of the Labor Market

Similar to Card et al. (2018) or Lamadon et al. (2022), suppose that L workers choose employers among a market of N firms, and normalize the aggregate mass of workers $L = 1$ for simplicity. Worker i 's utility from working at firm j is increasing in firm-specific amenities

³²Profits per worker type i are $(\Pi_{j,t}^i - 1) \times W_{j,t}$ where $\Pi_{j,t}^i$ is firm j net inverse wage markdown for type i , as defined in section 4.

$a_{j,t}$, the log of the wage offer $w_{j,t}$, and an unobservable taste shock $\epsilon_{i,j,t}$.

$$u_{i,j,t} = \varepsilon \log(w_{j,t}) + a_{j,t} + \epsilon_{i,j,t} \quad (\text{IA.32})$$

Assuming the taste shocks follow a type I extreme value distribution, then standard results from [McFadden \(1973\)](#) imply the firm-specific labor supply curve:

$$L(w_{j,t}, a_{j,t}) = \lambda_t^{-1} \exp(a_{j,t}) w_{j,t}^\varepsilon \quad (\text{IA.33})$$

Like [Card et al. \(2018\)](#) and [Lamadon et al. \(2022\)](#), here I assume each firm views itself as atomistic in the market, and so they take the constant $\lambda_t = \left(\sum_{j'=1}^N \exp(a_{j',t}) w_{j',t}^\varepsilon \right)$ as given. The parameter ε gives the firm-specific supply elasticity. Let $A_{j,t} = \lambda_t^{-1} \exp(a_{j,t})$ denote the level of the supply curve.

Firms choose the wage offer $w_{j,t}$ to maximize the following:

$$V_{j,t} = \max_{w_{j,t}} Z_{j,t} (L_{j,t})^{1-\alpha} - w_{j,t} L_{j,t} \quad (\text{IA.34})$$

subject to the functional form of the labor supply curve ([IA.33](#)). Denote the wage markdown by $\mu \equiv \frac{\varepsilon}{\varepsilon+1}$ and define constant $c \equiv \frac{1}{1+\varepsilon\alpha} \log(\mu(1-\alpha))$. Solving ([IA.34](#)) yields the expressions for the log optimal employment, wage, and firm value:

$$\log(L_{j,t}) = \varepsilon c + \frac{\varepsilon}{1+\varepsilon\alpha} \log(Z_{j,t}) + \frac{1}{1+\varepsilon\alpha} \log(A_{j,t}) \quad (\text{IA.35})$$

$$\log(w_{j,t}) = c + \frac{1}{1+\varepsilon\alpha} \log(Z_{j,t}) - \frac{\alpha}{1+\varepsilon\alpha} \log(A_{j,t}) \quad (\text{IA.36})$$

$$\log(V_{j,t}) = \log(1 - (1-\alpha)\mu) + (1-\alpha)\varepsilon c + \frac{1+\varepsilon}{(1+\varepsilon\alpha)} \log(Z_{j,t}) + \frac{1-\alpha}{1+\varepsilon\alpha} \log(A_{j,t}) \quad (\text{IA.37})$$

For simplicity I look at the bias from running a regression of the levels of log employment/wages on log firm value. My empirical strategy of running the regression of stock returns on the growth rates in wages and employment is essentially the same except in differences instead of levels. The parameter estimate obtained from regressing log employment and wages on firm value and taking the ratio of the two coefficients is given by

$$\hat{\varepsilon}_{IV}^{\log V} = \frac{\varepsilon(1+\varepsilon)\sigma_z^2 + (1-\alpha)\sigma_a^2 + \varepsilon(1-\alpha)\sigma_{az} + (1+\varepsilon)\sigma_{az}}{(1+\varepsilon)\sigma_z^2 + (1-\alpha)\sigma_{az} - \alpha(1+\varepsilon)\sigma_{az} - \alpha(1-\alpha)\sigma_a^2} \quad (\text{IA.38})$$

Here σ_z^2 is the variance of $\log(Z_{j,t})$; σ_{az} is the covariance of $\log(Z_{j,t})$ and $\log(A_{j,t})$; and, σ_a^2 is the variance of $\log(A_{j,t})$. When there are no labor supply shocks, so that $\sigma_a^2 = \sigma_{az} = 0$, equation ([IA.38](#)) collapses to the true supply elasticity ε .

For ([IA.38](#)) to represent an upper bound on the supply elasticity—which implies a

conservative estimate of the magnitude of wage markdowns—we must have

$$\rho_{az}\sigma_z > -\frac{(1-\alpha)}{1+\varepsilon}\sigma_a \quad (\text{IA.39})$$

where ρ_{az} is the correlation of $\log(A_{j,t})$ and $\log(Z_{j,t})$. The intuition behind (IA.39) is simple. Increases in the level of the supply curve $\log(A_{j,t})$, (increases in “amenities”) allow firms to hire more workers at a given wage, which reduces wages, increases employment, and increases firm value, all else held constant. This tends to bias the wage response downward and the employment response upward, leading to an upward biased elasticity estimate. However, if $\log(A_{j,t})$ is sufficiently negatively correlated with firm productivity, then increases in $\log(A_{j,t})$ reduce firm value, and the elasticity estimate becomes downward biased.

The most obvious candidate for reversing the inequality (IA.39) is market-specific productivity shocks. This is because $A_{j,t}$ is decreasing in the market-wide wage index $\lambda_t = \left(\sum_{j'=1}^N \exp(a_{j',t})w_{j',t}^\varepsilon\right)$:

$$A_{j,t} = \lambda_t^{-1} \exp(a_{j,t}) = \left(\sum_{j'=1}^N \exp(a_{j',t})w_{j',t}^\varepsilon\right)^{-1} \exp(a_{j,t}) \quad (\text{IA.40})$$

Suppose that firm-specific productivity $Z_{j,t}$ is given by an aggregate market component and an idiosyncratic component: $Z_{j,t} = \tilde{X}_t X_{j,t}$. Equation (IA.36) then implies that the wage offer can be expressed as

$$w_{j,t} = C_{j,t} \tilde{X}_t^{\frac{1}{1+\alpha\varepsilon}} \lambda_t^{\frac{\alpha}{1+\alpha\varepsilon}} \quad (\text{IA.41})$$

where the firm-specific term $C_{j,t}$ depends on $\exp(a_{j,t})$, $X_{j,t}$, and model parameters. Define $\tilde{C}_{j,t} = \exp(a_{j,t})C_{j,t}^\varepsilon$. Then

$$\lambda_t^{\frac{1}{1+\alpha\varepsilon}} = \left(\sum_{j'} \tilde{C}_{j',t}\right) \tilde{X}_t^{\frac{\varepsilon}{1+\alpha\varepsilon}} \equiv \tilde{C}_t \tilde{X}_t^{\frac{\varepsilon}{1+\alpha\varepsilon}} \quad (\text{IA.42})$$

or

$$\lambda_t = \tilde{C}_t^{1+\alpha\varepsilon} \tilde{X}_t^\varepsilon \quad (\text{IA.43})$$

Hence $\log(A_{j,t}) = a_{j,t} - (1+\alpha\varepsilon) \log(\tilde{C}_t) - \varepsilon \log(\tilde{X}_t)$ is decreasing in $\log(\tilde{X}_t)$. If enough of the variance in $\log(A_{j,t})$ is driven by $\log(\tilde{X}_t)$ this can reverse the inequality in (IA.39), causing downward biased supply elasticities. Aggregate productivity shocks do have meaningful variance and firm-specific amenities may be much more slow moving, so this case is feasible empirically. Since I assume firms are “small” in the market, the correlation of $\log(\tilde{C}_t)$ and $\log(C_{j,t})$ is approximately zero and so I don’t consider the impact of this term. The $\log(\tilde{C}_t)$ term nets out with market-wide fixed effects along with $\log(\tilde{X}_t)$. In summary, this means

that without market-specific controls I may overestimate the importance of labor market power.

After netting out market-specific productivity shocks (via market-by-year fixed effects or other controls) the condition (IA.39) becomes highly plausible. When market shocks are netted out, (IA.39) says that workers' perceptions of firm amenities should not *decrease* by too much when idiosyncratic productivity *improves*, or that firms should not cut their amenities by a lot on average when they experience a positive productivity shock. If there is any correlation at all, a more likely scenario is that firm amenities also improve when their productivity goes up, which guarantees that (IA.39) holds, implying my estimates would be conservative if biased at all. Related restrictions on unobservable firm-specific amenities are imposed in papers that estimate passthrough parameters of firm-specific shocks to workers, including Lamadon et al. (2022), Kline et al. (2019), and Garin and Silverio (2020).³³ The same issue can arise with quasi-experimental variation in firm revenue productivity, such as the construction-lottery designs of Lamadon et al. (2022) and Kroft et al. (2025), because firms could adjust amenities in response to the exogenous shock.

Now, compare the IV estimate of the elasticity obtained in equation (IA.38) with the OLS estimate from regressing log employment on log wages:

$$\hat{\varepsilon}_{OLS} = \frac{\varepsilon\sigma_z^2 - \alpha\sigma_a^2 + (1 - \varepsilon\alpha)\sigma_{az}}{\sigma_z^2 + \alpha^2\sigma_a^2 - 2\alpha\sigma_{az}} \quad (\text{IA.44})$$

In contrast to the IV estimate, the OLS elasticity estimate in (IA.44) can be biased up or down depending on parameters, even after invoking assumption (IA.39). To the extent it is better to know that wage markdown estimates are conservative, if biased at all, then the IV estimates are clearly preferable to the OLS. In practice, there is also likely some measurement error in firms' wages, for example due to imperfections in the linkages between firm wage information in the LEHD and financial information in Compustat, as well as incomplete coverage of firms' establishments in the LEHD. This would attenuate OLS estimates of elasticities towards zero and yield overly large wage markdowns, while an instrumental variables estimate corrects for this measurement error.

IA.4 Elasticity Estimates with Noise in Firm Value

Consider the following modification to the firm's static problem (1):

$$V(\Theta, \Gamma, \Lambda) = G(\Lambda) + \max_L [F(L; \Theta) - W(L; \Gamma)L] \quad (\text{IA.45})$$

³³Papers in the asset pricing literature which incorporate labor market frictions, such as Kuehn et al. (2017), Belo et al. (2014), and Donangelo (2014), also assume that labor supply shocks are determined at the market level, implying that firm-specific stock returns net of within-market effects yield valid identifying variation.

Here $G(\Lambda)$ affects firm value, but not labor demand. Suppose there is a shock both to some $\lambda \in \Lambda$ and some $\theta \in \Theta$. Define the estimate of the supply elasticity:

$$\left[\hat{\epsilon}_{V,\theta,\gamma} = \left(\frac{dL}{dV} \right) / \left(\frac{dW}{dV} \right) \right] \times \frac{W}{L} \quad (\text{IA.46})$$

Define $\Phi(\Theta, \Gamma) = \max_L [F(L; \Theta) - W(L; \Gamma)L]$. Taking total derivatives, we have

$$dV = \frac{\partial G}{\partial \lambda} d\lambda + \frac{\partial \Phi}{\partial \theta} d\theta \quad (\text{IA.47})$$

$$dL = \frac{\partial L}{\partial \theta} d\theta \quad (\text{IA.48})$$

$$dW = \frac{\partial W}{\partial L} \frac{\partial L}{\partial \theta} d\theta \quad (\text{IA.49})$$

Where in an abuse of notation L and W correspond to their equilibrium values at the firm's optimum. Then

$$\hat{\epsilon}_{V,\theta,\gamma} = \left(\frac{dL}{dV} \right) / \left(\frac{dW}{dV} \right) \times \frac{W}{L} = \left(\frac{\partial L}{\partial \theta} d\theta \right) \times \left(\frac{\partial W}{\partial L} \frac{\partial L}{\partial \theta} d\theta \right)^{-1} \times \frac{W}{L} = \frac{\partial L}{\partial W} \frac{W}{L} = \epsilon \quad (\text{IA.50})$$

Since the dV terms cancel out, the existence of $G(\Lambda)$ has no effect on the derived supply elasticity estimate.

IA.5 Wage Markdowns Under a Multilateral Bargaining Framework

In section 4, I show that minor extensions on a wage-posting, monopsonistic setup, borrowing a labor supply functional form from [Card et al. \(2018\)](#) and [Kline et al. \(2019\)](#), can accurately describe many important empirical regularities, including variable supply elasticities by firm productivity. My illustrative example in (8) also used this same labor supply functional form and mode of firm-specific wage and employment determination.

Besides monopsonistic wage-posting models, other models of imperfect competition in labor markets imply that firm-specific shocks to marginal productivity should affect wages and employment. Here I explore the interpretation of firm-specific shocks on labor demand when wages are instead determined by multilateral bargaining in multi-worker firms, as in [Stole and Zwiebel \(1996\)](#) and [Acemoglu and Hawkins \(2014\)](#), and subject to convex hiring costs, as in [Garin and Silverio \(2020\)](#). The multilateral bargaining setup yields a “quasi-elasticity” (defined as the ratio of equilibrium log employment to log wage) of labor supply that is very closely related to the structural labor supply elasticity in the wage posting frameworks in ([Card et al., 2018](#); [Kline et al., 2019](#)).

Inferred wage markdowns from applying equation (4) to estimate the wage markdown are quantitatively close to the case when firms instead engage in multilateral bargaining (and are likely to be conservative for reasonable parameterizations of convexity in the hiring cost function). In the paper, I take the wage posting perspective because it is particularly useful for modeling and interpreting the heterogeneous labor demand responses to firm-specific shocks in terms of wage markdowns; it also has a straightforward extension to a setting with dynamic labor demand, which I discuss in section 4. To the extent that such bargaining motives also operate in labor markets, the below discussion makes clear that the close relationship between the two frameworks ensures that quantitative estimates for wage markdowns are not likely to be substantively biased upward by assuming wage posting.

I now explore how wage markdowns behave in a multilateral bargaining framework (Stole and Zwiebel, 1996; Acemoglu and Hawkins, 2014) where multi-worker firms hire subject to a convex hiring cost. This is closely related to the setup explored in the appendix of Garin and Silverio (2020). My focus on publicly traded firms, which have hundreds or thousands of employees, combined with the weak role of collective bargaining in the US labor market, together suggest that this is the best bargaining framework to benchmark my wage posting baseline against (as opposed to models of bargaining with single-worker firms or collective bargaining, for example).

The firm has a revenue function $F(L)$ with $F(L) = ZL^{1-\alpha}$, so that $\alpha > 0$ implies diminishing returns to scale, and pays a convex hiring cost $c(L) = \frac{\kappa}{\theta+1}L^{\theta+1}$ to hire L workers. The parameter Z gives firm total factor productivity. Wages are determined by the multilateral bargaining solution of (Stole and Zwiebel, 1996; Acemoglu and Hawkins, 2014), which generalizes Nash bargaining to large firms with diminishing marginal productivity. The firm solves:

$$\max_L F(L) - c(L) - W(L)L \quad (\text{IA.51})$$

Let η denote the workers' bargaining power. As in (Stole and Zwiebel, 1996; Acemoglu and Hawkins, 2014), the solution to the multilateral bargaining process yields the agreed upon wage

$$W(L) = (1 - \eta)b + \gamma\eta F_L(L) \quad (\text{IA.52})$$

with $\gamma \equiv \frac{1}{1-\alpha\eta}$ and b the workers' outside option. Because the firm has decreasing returns to scale, the multilateral bargaining implies lower marginal revenue product as L increases, which means the wage offer is smaller for a higher L for a fixed level of productivity Z . I assume Z is sufficiently high that the firm deems it worthwhile to hire a positive amount of workers.

The first-order conditions are:

$$F_L(L) = c_L(L) + W_L(L)L + W(L) \quad (\text{IA.53})$$

I solve for the expression for the quasi labor supply elasticity

$$\frac{d \log L^*}{d \log W^*} \quad (\text{IA.54})$$

where L^* is the solution to the firm's problem (IA.51) and $W^* = W(L^*)$. I do so by first solving for the inverse labor supply elasticity through a shock to Z :

$$\frac{dW^*}{dL^*} \frac{L^*}{W^*} = \left(\frac{\partial W^*}{\partial Z} \frac{\partial Z}{\partial L^*} + \frac{\partial W^*}{\partial L^*} \right) \frac{L^*}{W^*} \quad (\text{IA.55})$$

Going forward I remove the superscript $*$ from my notation, with the understanding that L and W refer to the equilibrium optimal employment and wages at the firm's solution to (IA.51). The components of (IA.53) are given by

$$F_L = Z(1 - \alpha)L^{-\alpha} \quad (\text{IA.56})$$

$$c_L = kL^\theta \quad (\text{IA.57})$$

$$W(L) = (1 - \eta)b + \gamma\eta Z(1 - \alpha)L^{-\alpha} \quad (\text{IA.58})$$

$$W_L L = -\alpha\gamma\eta Z(1 - \alpha)L^{-\alpha} \quad (\text{IA.59})$$

Collecting terms for (IA.53):

$$(1 - \alpha)ZL^{-\alpha} = kL^\theta - \alpha\gamma\eta(1 - \alpha)ZL^{-\alpha} + (1 - \eta)b + \gamma\eta(1 - \alpha)ZL^{-\alpha} \quad (\text{IA.60})$$

Applying the implicit function theorem to (IA.60) yields the expression for $\frac{\partial Z}{\partial L}$

$$\frac{\partial Z}{\partial L} = \frac{\theta k L^{\theta+\alpha} + \alpha(1 - \alpha)Z - \alpha\gamma\eta(1 - \alpha)Z + \alpha^2(1 - \alpha)\eta\gamma Z}{[(1 - \alpha) - \gamma\eta(1 - \alpha) + \alpha\gamma\eta(1 - \alpha)] L} \quad (\text{IA.61})$$

Next, differentiating (IA.58) with respect to Z :

$$\frac{\partial W}{\partial Z} = \eta\gamma(1 - \alpha)L^{-\alpha} \quad (\text{IA.62})$$

Finally, differentiating (IA.58) with respect to L gives

$$\frac{\partial W}{\partial L} = -\alpha\eta\gamma(1 - \alpha)L^{-\alpha-1} \quad (\text{IA.63})$$

Plugging in (IA.61), (IA.62), and (IA.63) to (IA.55), applying the definition of γ , and algebraic manipulation results in the inverse quasi-elasticity:

$$\frac{d \log W}{d \log L} = \frac{\theta k L^\theta}{\frac{1-\eta}{\eta} W} \quad (\text{IA.64})$$

By manipulating the first order conditions, we can substitute $kL^\theta \equiv c_L(L) = \frac{1-\eta}{\eta}(W-b)$. Substituting and taking the reciprocal of (IA.64) yields the desired quasi-elasticity of labor supply

$$\frac{d \log L}{d \log W} = \frac{W}{\theta(W-b)} \equiv \tilde{\varepsilon} \quad (\text{IA.65})$$

Compare this expression to the supply elasticity generated from a wage posting model where the firm-specific labor supply curve is proportional to $(W-b)^\beta$, as in Card et al. (2018), Kline et al. (2019), and in my framework in section 4 of the main text:

$$\hat{\varepsilon}_{\text{Wage posting}} \equiv \frac{\beta W}{(W-b)} \quad (\text{IA.66})$$

The connection between the two is readily apparent: the labor supply curve slope parameter β in the wage posting, monopsonistic framework and $(1/\theta)$, the reciprocal convexity of the hiring cost function in the multilateral bargaining framework, are observationally equivalent to one another, and both models deliver decreasing labor supply elasticity estimates as labor productivity $F/L \propto F_L$ increases.

Now, consider the wage markdown that would be inferred by estimating the quasi supply elasticity and assuming that the firm behaves as a wage posting monopsonist when multilateral bargaining is the true determinant of wages:

$$\text{Quasi Markdown} = \frac{\tilde{\varepsilon}}{1 + \tilde{\varepsilon}} = \frac{W}{(1 - (1 + \theta)\eta)b + (1 + \theta)\eta\gamma F_L(L)} \quad (\text{IA.67})$$

Compare this to the true wage markdown under multilateral bargaining, the fraction of marginal product that workers receive in wages:

$$\text{True Markdown} = \frac{W}{F_L(L)} \quad (\text{IA.68})$$

The true markdown (IA.68) under multilateral bargaining also implies that wage markdowns are increasing in labor productivity, similar to my model in section 4. The relative bias ratio in wage markdowns from assuming wage posting monopsony instead of multilateral bargaining is

$$\frac{\text{True Markdown}}{\text{Quasi Markdown}} = \frac{(1 - (1 + \theta)\eta)b + (1 + \theta)\eta\gamma F_L(L)}{F_L(L)} \quad (\text{IA.69})$$

To understand the size of the bias, consider the special case $1 + \theta = 1/\eta$. Then (IA.69) collapses to $\gamma = \frac{1}{1-\alpha\eta}$. Because structural estimates of η are small, reasonable calibrations of α and η imply that the constant γ is close to 1. For example, Kuehn et al. (2017) estimate $\eta = 0.115$ and Hagedorn and Manovskii (2008) estimate $\eta = 0.052$. Taking $\eta = 0.115$ from Kuehn et al. (2017), which implies a larger γ and is hence more conservative, and following their calibration of labor returns to scale $1 - \alpha = 0.75$, this gives $\gamma \approx 1.029$, implying that markdowns inferred by incorrectly assuming a monopsonistic wage determination instead of multilateral bargaining would cause one to slightly overestimate the size of wage markdowns from marginal product (to the tune of about 3% in this example) when $1 + \theta = 1/\eta$.

Note however that imposing $1 + \theta = 1/\eta$ entails a very high amount of convexity in hiring costs: for example, using $1/\eta = 1/0.115$ implies an exponent of about 8.7 in the hiring cost function, while it is common to specify quadratic vacancy costs in bargaining models (Acemoglu and Hawkins (2014), for instance). Thus I also consider the size of the relative bias when the convexity is less than the reciprocal of the bargain weight: $1 + \theta < 1/\eta$. To make progress, first observe that (IA.60) can be rearranged as follows:

$$\gamma F_L(L) = \frac{kL^\theta}{1 - \eta} + b \quad (\text{IA.70})$$

which implies $\gamma F_L(L) > b$.³⁴ This further implies that the numerator in (IA.69) is strictly increasing in $(1 + \theta)$. Hence the ratio must be strictly decreasing as convexity θ decreases, and in turn smaller than γ when $1 + \theta < 1/\eta$. The ratio (IA.69) can be well below 1 for a wide range of feasible values for $1 + \theta$. Therefore any possible upward bias in the extent of market power from incorrectly assuming a wage posting, monopsonistic labor market instead of multilateral bargaining must be small in practice, and for economically reasonable parameterizations of hiring cost convexity, the monopsony assumption yields conservative wage markdowns.

IA.6 Supply Elasticity Robustness Checks: Empirical Specifications and Data

IA.6.1 Different Controls for Market Level Shocks

Failing to account for common market shocks could bias my supply elasticity estimates downward, leading to an upward bias in the extent of labor market power. In Table IA.3 I present the elasticity estimates implied by estimating my baseline specification (9) with homogeneous supply elasticities (panel A), and also allowing them to vary based off above- or

³⁴Note that this rearrangement of the FOC also illustrates that equilibrium marginal revenue product is constant if $\theta = 0$. Hence if hiring costs are linear then marginal product always adjusts such that shocks to Z do not move wages.

below-median labor productivity. Across the columns I explore alternative specifications with different controls for common market shocks. My estimate of the average elasticity—obtained by taking the ratio of the employment and wage responses—in the baseline specification (column 1) is about 2.5. In columns 2-9 of Table [IA.3](#) I include different variations of controls for common market shocks.

I find that each set of controls for common shocks yields estimates that are quantitatively very close, and well within the 95% confidence interval of my baseline estimate. From this exercise I conclude that the main specification in column 1 of Table [IA.3](#) does a reasonably good job of capturing the relevant variation in common market shocks that could bias my estimates downward. Consequently I focus on this baseline specification, noting that any quantitative changes from using a different specification would be minor.

In column 2 of Table [IA.3](#) I drop the industry-by-year fixed effects. Consistent with the discussion above and the simple model in internet appendix [IA.3](#), this reduces my supply elasticity estimate, but only slightly (2.525 in column 1 versus 2.397 in column 2). In the third column of Table [IA.3](#) control for the weighted average employment and wage growth of a firms' labor market competitors instead of industry-by-year fixed effects. To create this measure I weight the employment or wage growth by the number of workers a competitor has hired from the given firm, divided by the total number of workers that have been hired away by other firms in the sample. I similarly do this for the number of workers a given firm hires from a candidate firm. I then take the average of the inflow- and outflow-weighted measures. Controlling for these variables also do not move the estimated elasticity substantially.

In the fourth column I control for empirical labor market-by-year fixed effects instead of the industry-by-year fixed effects from column 1. I obtain these empirical labor market boundaries by performing k-means clustering on the flows of workers between firms; I explain the method in more detail in section [IA.8](#) of the internet appendix and show that it does a good job of capturing variation in worker flows between firms, as well as in wages and employment levels and growth rates. I choose $k = 10$ labor market clusters per year as my baseline.³⁵ See appendix Table [IA.13](#) for more details on the comparative performance of different definitions of labor market boundaries. In column 4 the elasticity is about 2.54 when controlling for these empirical labor markets. In column 5 I use $k = 20$ labor market clusters per year instead of $k = 10$, which yields an elasticity estimate of 2.57

In column 6 I control for both industry-by-year and labor market-by-year fixed effects, and in column 7 do the same using the $k = 20$ labor market version. In column 8 I also add back in the competitor employment and wage growth controls on top of the $k = 10$ version of labor market fixed effects. Finally, in column 9 I add in local labor market controls,

³⁵Using a different clustering method, [Nimczik \(2020\)](#) finds that $k = 9$ empirical labor markets does a good job of capturing empirical market boundaries in Austria.

including the employment-weighted average changes in local labor market concentration and unemployment rates in the commuting zones across all the firm’s establishments, as well as the average stock returns of other public firms who are headquartered in the same commuting zone; I additionally include the competitor average stock return, which is the average stock return over same labor market competitors as the wage growth controls. Across all these permutations the estimates are remarkably stable, varying between 2.40 and 2.61, and all falling well within the 95% confidence interval of my baseline estimate of 2.525 (with a standard error of 0.305). This exercise demonstrates that my estimates are not sensitive to the type of market-level controls that I include, and so I take my industry-by-year fixed effect specification to be a reasonable parsimonious baseline specification.

Examining panel B with labor productivity heterogeneity, the takeaway is the same: not only are homogeneous supply elasticities highly stable, altering the set of controls has no meaningful effect on the supply elasticity gap between high- and low labor productivity firms. By and large, this exercise illustrates that the manner in which I account for market-level shocks does not play a meaningful role in my quantitative findings.

IA.6.2 Within Local Labor Market Sorts

The main analysis in the paper analyzes employment at the Compustat firm level. While this aligns with the level at which we can observe firm-level financial information, it doesn’t allow me to account directly for local labor market shocks. To the extent that the same firm may have a presence in multiple local labor markets, analyzing employment and wage responses directly at the local level allows for more granular labor market fixed effects. However, the LEHD does have some limitations in this respect; it reports individual wages at the state tax entity identifier level (called the “SEIN”). For multi-establishment firms, the SEIN typically groups establishments together, although the boundaries of the SEIN vary somewhat ([Vilhuber, 2018](#)). Still, as ([Vilhuber, 2018](#)) discusses, much of LEHD employment comes from single-establishment entities, and Compustat firms often comprise of many SEINs. The LEHD also provides a county and industry assignment for each SEIN (where the county of highest employment is used when SEIN establishments span multiple counties).

Using this local granularity, I map counties to Census commuting zones, and I follow [Berger et al. \(2022\)](#) in defining local labor markets as the set of SEINs operating in the same industry (3-digit NAICS) and commuting zone. I follow all the exact same procedures for wage cleaning, constructing the sample, and variable construction as described in Appendix A, except I now define firms at the SEIN level. Because Compustat employment is not available for SEINs, I define employment as the total number of workers who are employed at the SEIN as their primary employer in the last quarter of the year, and who satisfy the sample requirements described in Appendix A.

Let j denote an SEIN, and $\alpha_{M(j),t}$ local labor market $\times$ year fixed effects. I sort firms

based on two characteristics. First, similar to [Berger et al. \(2022\)](#), I sort on the firm’s local labor market wage bill share. Since my LEHD sample only includes public firms, I compute the total wage bill of the local market by summing payrolls for all firms in Longitudinal Business Database (LBD) who are operating in the same local market, and then I define the SEIN’s market wage bill share as the ratio of SEIN LEHD wage bill over total LBD market wage bill. I sort SEIN-years into 4 quartiles based off local wage bill share. Next, motivated by the local spatial sorting model of [Lindenlaub et al. \(2024\)](#), I rank firms based on their (log) labor productivity advantage, relative to the average labor productivity of other firms operating in the same market. This measure has the advantage of being closer to my main specification, but also has the disadvantage that I can only compare with the labor productivity of other public firms in the same market. I also have to assume that all SEINs coming from the same Compustat firm have the same labor productivity as the overall corporation. I sort SEINs cross-sectionally into 4 quartiles based off this measure of local labor productivity advantage. Let $q(j, t)$ denote the quartile of SEIN j in year t based off either their local labor market share or local labor productivity advantage. I estimate the following specification via instrumental variables:

$$\log EMP_{j,t+1} - \log EMP_{j,t} = \alpha_{q(j,t)} + \alpha_{M(j),t} \sum_{q=1}^2 \varepsilon_q \times \mathbf{1}(q(j, t) = q) \times (\log W_{j,t+1} - \log W_{j,t}) + \Psi X_{j,t} + \epsilon_{j,t+1} \quad (\text{IA.71})$$

I again use firm stock returns as the instrument to estimate elasticities. Relative to my main analysis, the specification [\(IA.71\)](#) has the advantage of granular local market fixed effects $\alpha_{M(j),t}$. Unlike my main specification, it also leverages both between-firm and within-firm variation to identify elasticity parameters, as it allows an SEIN with a larger labor market competitive advantage to have a different supply elasticity compared with an SEIN coming from the same Compustat firm, but with a smaller labor market competitive advantage.

Estimates are reported in appendix Table [IA.4](#). As discussed in the main text, I find lower supply elasticities for firms with either a higher local labor market wage bill share, consistent with [Berger et al. \(2022\)](#), and also lower elasticities for firms with a larger local labor productivity advantage, consistent with [Lindenlaub et al. \(2024\)](#) and the main analysis in this paper. These estimates are qualitatively consistent with my main analysis, but because LEHD employment is somewhat less responsive to stock return shocks, I obtain generally lower labor supply elasticities using this alternative formulation, meaning my main analysis generates smaller markdowns than would be implied by these local sorts. Still, it is reassuring that this local variation generates patterns that are largely consistent with the rest of the paper, and also that the inability to control for more granular local shocks in my main specifications is not likely to bias my estimates. This is also consistent with the analysis in appendix Table [IA.3](#), which showed that elasticity estimates do not change when I explore

adding Compustat-firm level controls for average exposure to local market shocks.

IA.6.3 Alternative Labor Demand Shocks

In this section I provide more details on the alternative labor demand shifters explored in Table 4 of the main text.

In the top panel of Table 4 I use stock returns of firms' customers rather than the firms themselves as shock to labor demand. [Cohen and Frazzini \(2008\)](#) show that stock return shocks to customers eventually propagate to upstream to their suppliers as a demand shock, and consistent with this notion I find that the wages and employment at supplier firm both respond significantly positively to the stock returns to their customers. The identifying assumption is that these shocks to customers constitute a pure demand shock to the firm, and are orthogonal to firm-specific labor demand shocks after accounting for industry-by-year fixed effects. Customers' returns are likely much less affected by any idiosyncratic labor supply shocks of their suppliers than the suppliers' own returns.

In the next panel of Table 4 I use patent-induced shocks from [Kogan et al. \(2017\)](#) as a shifter of firm labor demand. [Kogan et al. \(2017\)](#), show that their measure predicts changes in firm productivity, employment, and sales, all consistent with a marginal revenue productivity shock. Patent values are estimated from stock price movements in a small window around patent grants, and capture information in price movements related to firm innovation. I follow [Kogan et al. \(2020\)](#) in looking at the response of valuable of patents from the year they are filed rather than granted, and use patenting in year t as a shock to labor demand from year t to $t + 1$. This specification is related in spirit to [Kline et al. \(2019\)](#), who estimate passthroughs of patent-induced shocks to worker and firm outcomes based on the timing of allowances of ex-ante "high-value" patents, based on predicted [Kogan et al. \(2017\)](#) patent values. Again I find an elasticity estimate that is close to my baseline.

I next construct firms' earnings announcement surprises following [Daniel, Hirshleifer, and Sun \(2019\)](#) and [Chan, Jegadeesh, and Lakonishok \(1996\)](#), by using stock returns in excess of the market return in a four day window around earnings announcements, which isolates periods of time when fundamental information about firm cashflows is revealed. Finally, following [Daniel et al. \(2019\)](#) and [Chan et al. \(1996\)](#) I construct earnings announcement surprises by calculating cumulative daily returns in excess of the market for the four-day period starting the day before an earnings announcement, and then average across all 4-day announcement returns over the July through June period.

The goal of this measure is to isolate a component of stock price movements that are highly likely to be related to information about firm productivity and unrelated to information about labor supply. Stock price movements around earnings announcement are driven primarily by firms announcing unexpectedly high or low earnings; thus restricting to these small windows is more likely to capture price movements related to information about product markets, and

hence firm labor demand. To the extent that earnings announcement reactions are driven by information about firm revenue productivity and not news about labor supply (which likely doesn't move at such a high frequency), this measure isolates valid identifying variation in stock returns. Though the measures are different, this follows a similar reasoning as my use of [Kogan et al. \(2017\)](#) patent induced shocks to the firm—both measures isolate movements in prices due to information revealed to the market that is highly related to firm revenue productivity, and hence instruments for shifts in labor demand.

Each different labor demand shifter in Table 4 yields elasticity estimates that are very close quantitatively and statistically indistinguishable from my baseline estimates, and each demonstrates a decreasing slope in supply elasticities as labor productivity improves. The commonality in qualitative and quantitative takeaways across different set of instruments suggests that any bias from any confounding firm-specific labor supply shocks is not likely to have large a quantitative effect on my stock return-based supply elasticity estimates.

IA.6.4 Controlling for Observable Labor Supply Shocks (Union Elections and Changes in Non-Compete Enforceability)

While these findings are useful in establishing the plausibility of my baseline supply elasticity estimates, it would still be helpful to control for potential firm-specific labor supply shocks, if they can be made observable. Although one can never definitively say that every firm-specific labor supply shock has been accounted for, I now include specifications where I control separately for two labor market shocks that have featured prominently in prior literature: union elections and changes in non-compete contract enforceability. Results are in appendix Table IA.11, which shows that bias from excluding these more salient observable firm-specific labor supply shocks is not important quantitatively. This lends credence to my argument that this sort of confounding variation is not likely to substantially bias my estimates in general.

I control for unionizations using data on union elections from the National Labor Relations Board and matched to Compustat records by [Knepper \(2020\)](#).³⁶ Changes in non-compete enforceability come from the lists of changes compiled by [Ewens and Marx \(2017\)](#) and [Jeffers \(2019\)](#). Following [Ewens and Marx \(2017\)](#) I assign non-compete changes by firms' headquarters.

Union elections may be one of the single best candidates for the type of firm-specific labor supply shock that could bias my estimates. For example, [Lee and Mas \(2012\)](#) find that union elections wins induce significant negative stock return responses; I verify the same result in the last row of Panel A of Table IA.11. However, even among the firms who have experienced large union elections, the unionization event accounts for a small amount of variation in firm-specific stock returns, so that controlling for unionizations leads to negligible changes in

³⁶Thanks to Matthew Knepper for generously sharing his data.

my estimates. Meanwhile, [Jeffers \(2019\)](#) argues that firms use non-compete successfully to diminish the mobility of their skilled workers, and so changes in non-compete enforceability could also in theory constitute a labor supply shock that bias my estimates. .

In Panel A of appendix Table [IA.11](#) I re-estimate supply elasticities based off (9) for firms who ever experienced a union election during my sample; the first column is from the baseline specification without controlling for unionizations, and the others include different sets of unionization controls. Consistent with [Lee and Mas \(2012\)](#), I find a significantly negative stock return response in the year a union wins an election. Despite this fact, allowing for this supply shock to be observable has minimal effect on my estimated elasticities, despite the fact that I restrict the sample to only firms who have experienced a sufficiently large union election.

Following [Knepper \(2020\)](#), I focus on firms experiencing union elections where at least 20 employees voted in the election. In order to give maximal explanatory power to the union elections, I restrict the sample to just the set of firms identified at some point to have experienced a sufficiently large union election.

In the second column of Table [IA.11](#) I include dummies for whether a union election win occurs in year t , $t + 1$, or $t - 1$ (as well as dummies for whether any election is occurs); in the third column I ascribe all variation in stock returns in a union election year to the election by adding interactions of stock returns with the full set of union election dummies in the second column. In all cases the elasticity is quantitatively close to the baseline elasticity and statistically indistinguishable, even in the extremely conservative final column where my controls ascribe all variation in stock returns during the 3 years surrounding surrounding a union election to the election itself.

Non-compete agreements make it more difficult to move to a competing employer; [Jeffers \(2019\)](#) shows that they are quite common and especially prevalent among skilled workers. Changes in non-compete enforceability are therefore another good candidate for firm-specific labor supply shocks that could bias my estimates. Following [Ewens and Marx \(2017\)](#) I assign non-compete changes by firms' headquarters. Panel B of Table [IA.11](#) shows the resulting supply elasticities when controlling for non-compete changes. In the first column I include indicators for whether a non-compete increases or decreases in enforceability in years t , $t + 1$, or $t - 1$. In the next columns I add interactions of all non-compete dummies with the stock return in that year. As was the case with union elections, after controlling for non-compete changes, supply elasticity estimates are quantitatively very close and statistically indistinguishable, again even for the most conservative specification in the final column where the controls allow all variation in returns in the 3 years surrounding the non-compete law change event to be due to the law change itself.

IA.7 Additional Robustness Checks for Labor Productivity-Sorted Supply Elasticity Estimates

IA.7.1 Alternative Sorting Variables and Measures of Firm Wages/Employment

Since production methods and labor markets vary from industry to industry, one concern with my Table 3 could be reliance on unconditional labor productivity sorts that use between rather than within-industry variation. In Table IA.8 I instead sort firms into productivity quartiles within their 2-digit NAICS industry and re-estimate (13) for employment and wages to back out supply elasticities for workers of all skill levels. All my basic findings from Table 3 are unchanged for the within-industry sorts in IA.8, and all the elasticity point estimates are very similar.

In appendix Table IA.9 I address a few more potential concerns with elasticity estimates for firms sorted on productivity. One possibility is that wage responses are larger in the short-run for productive firms because they have more immediate flexibility in adjusting their wages, and so the monotonically decreasing pattern in elasticities sorted on productivity could be driven by the horizon. For example, unproductive firms may be more constrained in their ability to adjust wages in the short horizon. To address this, I re-estimate elasticities from (13) for the 3-year horizon. Specifically, I replace stock returns and employment/wage growth (as well as the control for contemporaneous changes in firm skill composition) with their 3-year equivalents:

$$\log(Y_{j,t+3}) - \log(Y_{j,t}) = \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} + \sum_{q=1}^4 \mathbf{1}(q(i,t) = q) \times \beta_q \text{Stock Ret}_{j,t \rightarrow t+3} + \Psi X_{j,t} + \epsilon_{j,t} \quad (\text{IA.72})$$

Panel A of Table IA.9 shows that productive firms similarly have much lower supply elasticities at this horizon. Thus the cross-sectional sorting in elasticities is not merely driven by the time horizon. Elasticities in general are also a little higher for this horizon, implying more elastic labor supply over the long run.

In Panel B of appendix Table IA.9 I re-estimate (13) with changes in the time-varying firm wage fixed effects from (A.5), instead of the log average firm wage. This is the component of the log wage that is entirely firm-specific and is paid to all workers, regardless of their skill, incumbent status, or worker-firm match quality; similar to the elasticity estimates for low- and middle-skill workers, this wage measure is less likely to be attached to firm- or worker skill-specific tendencies towards higher equity-based compensation. Again the most productive firms face by far the lowest supply elasticities. These elasticity estimates are not surprisingly a little higher than in Table 3, because can't take into account that the most highly paid workers make up a large share of the overall firm average wage and also have the least elastic labor supply. In internet appendix section IA.7.2 I detail more checks

which suggest that equity-based compensation is not driving the sorting patterns in supply elasticities.

In panel C of appendix Table [IA.9](#) I sort on estimates of firm total factor productivity from [İmrohoroglu and Tüzel \(2014\)](#) instead of labor productivity. Findings similarly go through in this case, as sorting on TFP also generates a decreasing pattern in supply elasticities. In unreported results I find a strong negative relationship between estimated TFP and firm labor shares. Hence my findings aren't just driven by my choice of the log value-added per worker measure, but are instead driven by productivity advantages in general.

Finally, due to concerns some have raised regarding employment reported in Compustat ([Davis, Haltiwanger, Jarmin, Miranda, Foote, and Nagypal, 2006](#)), in panel D of Table [IA.9](#) I replace the Compustat-based employment with Longitudinal Business Database employment. I obtain LBD employment by aggregating reported employment across all LBD establishments linked to a given Compustat gvkey in that year.³⁷ I still find strongly monotonically decreasing supply elasticities across productivity types, but the LBD employment is not as responsive to stock return shocks, and so I get slightly lower elasticity point estimates. This suggests that using LBD employment-based supply elasticities would increase the quantitative magnitude of my findings, if anything.

IA.7.2 Examining Whether Stock-Based Compensation Drives Sorting Patterns

Another concern may be that my findings for productive firms are driven by the equity-based compensation of the most skilled workers, who are disproportionately employed at productive firms. For example, [Eisfeldt et al. \(2021\)](#) document a large rise in equity compensation over my sample period, which may be due to contracting issues unrelated to a wage posting, monopsonistic model of the labor market. However, such equity-based incentive pay is not necessarily incompatible with a monopsony framework, as it intrinsically ties the compensation of employees with their marginal revenue productivity, as would be implied by a monopsony model.

The LEHD includes all compensation that is immediately taxable, which includes equity based pay upon exercise; however, as [Eisfeldt et al. \(2021\)](#) argue, a non-trivial fraction of equity-based compensation may never show up in LEHD earnings because they are taxed as capital gains, so it's not immediately obvious which direction the elasticities would be biased. Note also that these concerns are not only prevalent for my stock return-based labor demand shock, but for any shock that is correlated with firm performance (in other words, essentially any potential labor demand shifter). That being said, Table [3](#) already alleviates these concerns in part, because productive firms face significantly lower supply elasticities

³⁷LBD employment is collected in March while Compustat employment is almost always reported in December, and so I use LBD figures from the March nearest to the Compustat December employment report date to compute employment growth. I find that this yields larger employment responses—and hence more conservative elasticity estimates—than taking the previous March observation or an average of the two.

even for the least skilled workers, whose compensation is far less tied to equity-based incentive pay.³⁸

Stock-based compensation is more likely to be exercised and hence reflected in LEHD earnings when the firm performs well, so any bias from this channel would tend to show up more prominently for positive stock returns. In appendix Table IA.10 I allow wage responses to be asymmetric for positive and negative excess returns for firms sorted on labor productivity by estimating the following:

$$\begin{aligned} \log(w_{j,t+1}) - \log(w_{j,t}) = & \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} \\ & + \sum_{q=1}^4 \mathbf{1}(q(i,t) = q) \times \left[\beta_{1,q} \text{Stock Ret}_{j,t \rightarrow t+1} \times \mathbf{1}(\text{Stock Ret}_{j,t \rightarrow t+1} \geq 0) \right. \\ & \left. + \beta_{2,q} \text{Stock Ret}_{j,t \rightarrow t+1} \times \mathbf{1}(\text{Stock Ret}_{j,t \rightarrow t+1} < 0) \right] + \Psi X_{j,t} + \epsilon_{j,t} \end{aligned} \quad (\text{IA.73})$$

where again $q(j,t)$ gives firm j 's time- t labor productivity quartile, $\alpha_{I(j),t}$ are industry-year fixed effects, and $X_{j,t}$ are the same set of controls as throughout the paper.

The monotonic pattern in stock-return/wage response maintains on both the upside and downside, and wage responses are nearly indistinguishable for productive firms, while they are a bit lower for unproductive firms on the downside. Eisfeldt et al. (2021) report that smaller firms tend to use more equity-based pay, and these cross-sectional patterns in asymmetric wage responses appear consistent with that fact. Since elasticity estimates are decreasing in the wage response to stock return shocks, this implies there may be a slight relative downward bias in elasticities for unproductive firms, if anything.

To further address the role of stock-based or other deferred compensation in driving my findings, I also construct a measure to capture ex-ante reliance on this type of compensation. While stock-based compensation and perks for skilled employees may not be perfectly reflected in LEHD-reported compensation, Eisfeldt et al. (2021) argue that a significant part of sales, general, and administrative expenses (SG&A) reflect such costs, and hence can be used to create an expense-based measure of stock/incentive pay. Accordingly, I take the ratio of SG&A expenses to total LEHD-implied firm labor expenses to proxy for how intensively a firm relies on this form of compensation. I sort firms each year into above- and below-median bins for the SG&A to LEHD labor expense ratio, and then I group firms into labor productivity quartiles within these two categories.³⁹ In appendix Figure IA.6 I re-estimate (13) for these two groups; I find the same sorting pattern by labor productivity and nearly the exact same elasticity point estimates within each group, strongly suggesting that differential reliance on

³⁸Eisfeldt et al. (2021) find that 97% of equity-based incentive pay accrues to the top 10% of workers in the manufacturing sector.

³⁹As discussed by Eisfeldt and Papanikolaou (2013), standards for reporting SG&A expenses vary by industry, so I compute these bins within 2-digit NAICS category.

this type of compensation is not a major driving force behind my findings; the consistent elasticity point estimates across the two groups also suggests that besides cross-sectional differences in elasticities, bias in overall elasticity magnitudes due to issues related to stock compensation is not likely to be a first-order concern.

IA.8 Alternative to Industries to Determine Labor Markets: K-Means Clustering to Estimate Empirical Labor Market Boundaries

In the main text I use 3-digit NAICS industry-by-year fixed effects in my estimation strategy to difference out the within-market shocks to firms. However, given that industries are related to the mode of production it is possible these do an inadequate job of capturing within *labor market* shocks. In this section I describe an alternative method for grouping firms into empirical labor market boundaries based on k-means clustering of flows of workers across firms. While I show below that this strategy does a good job at capturing empirical labor market boundaries, appendix Table IA.3 demonstrates that any of a number of different reasonable proxies to difference out common market shocks yields extremely similar estimates of supply elasticities. Thus the way I control for these shocks is not of first-order quantitative importance (see section IA.6 of the internet appendix for more details on the various specifications in that table). Consequently, I stick with the simple industry definition of market boundaries in my main estimates and report other estimates as a robustness check.

Let $R_{i,j,t}$ be the number of workers firm i has hired from firm j in year t , and $L_{i,j,t}$ the number of workers firm i loses to firm j in year t . I then construct a matrix A_t of flows across firms by assigning the the i, j th entry as follows:

$$A_{i,j,t} = \log \left(1 + \sum_{\tau=t-2}^t (R_{i,j,\tau} + L_{i,j,\tau}) \right) \quad (\text{IA.74})$$

I then perform k-means clustering on the columns of A to group firms into clusters that hire from one another, using 1 minus the cosine similarity of column vectors of A_t as the distance metric between vectors. The cosine distance between column $A_{j,t}$ and $A_{j',t}$ is defined as

$$1 - \frac{\sum_i A_{i,j,t} \times A_{i,j',t}}{\sqrt{\sum_i A_{i,j,t}^2} \sqrt{\sum_i A_{i,j',t}^2}}, \quad (\text{IA.75})$$

which is one minus the uncentered correlation between the vectors $A_{j,t}$ and $A_{j',t}$. Thus the distance metric accounts for the differences in the size of the two comparison firms. I take the log of the sum of inflows and outflows in order to downweight extremely large firms. Because k-means clustering has a random component, I perform the routine 5 times each year using

different starting points for the clusters, selecting the cluster assignment which explains the largest share of workers flows across firms for that year. Because the number of clusters must be pre-specified, I perform the routine for both $k = 10$ and $k = 20$ labor market clusters per year.

I find that the method explains labor empirical labor markets quite well. In appendix table [IA.13](#) I examine how much variation labor market-by-year fixed effects explain in the levels of and changes in log wage and employment as well as the fraction of firm-to-firm worker transitions occurring within the given labor market boundary. I compare the $k = 10$ and $k = 20$ clusters with 2-digit and 3-digit NAICS industries. The table shows that about half of worker flows occur within the $k = 10$ version of the empirical labor markets, while also explaining three-fifths of log wages and having very comparable explanatory power for wage and employment growth, and stock returns as the other empirical labor market boundaries. In the final column of the table I show that, while the empirical labor market boundaries do a good job of capturing empirical labor markets, they capture similar variation in stock returns, and employment/wage growth as my baseline 3-digit NAICS fixed effects. In particular, there is a very small change in explanatory power when both labor market-by-year and 3-digit NAICS-by-year fixed effects are included at the same time.

IA.9 Interpreting the Magnitude of Cash Flows Generated from Wage Markdowns

Given the size of wage markdowns as a fraction of operating income in Table [6](#), it is helpful to provide some context for interpretation and to compare the magnitude against other results.

First of all, note that the markdown shares of capital income do not represent the change in the value of the firm relative to the competitive equilibrium. Constructing a reasonable competitive counterfactual would require imposing far more structure on my model. Rather, I am asking the question “What is the total value of the gap between wages and marginal products relative to total capital income, *holding equilibrium quantities fixed*?” In a competitive equilibrium quantities as well as prices would adjust. Hence my counterfactual does not imply, for example, that the aggregate enterprise value of publicly-traded firms would be 20-25% lower if we imposed perfect competition.

How do my markdown estimates compare to markdown or markup estimates in the literature? [Crouzet and Eberly \(2021\)](#) point out that the [De Loecker et al. \(2020\)](#) price markup estimates imply rents from product market power were worth almost 40% of value added by the end of their sample period. Taking the Compustat aggregate labor share in Table [7](#) and the aggregate markdown estimates in Table [6](#), my estimates imply wage markdowns are about 15% of value-added in the 2003-2014 period, implying the [De Loecker](#)

et al. (2020) estimates are more than two and a half times larger by this metric. As far as wage markdowns, Yeh et al. (2022) jointly estimate wage markdowns and price markups for manufacturing firms using a production function approach. They find that the average firm pays about 65% of marginal product. Berger et al. (2022) find an aggregate markdown of about 72% of marginal product. Meanwhile, my estimates imply that firms pay about 76% of aggregate total marginal product, smaller than both these estimates. This is despite the fact that my sample zeroes in on large, publicly traded firms, which likely have more market power, while their sample includes smaller private firms. Labor market power is not important for firms' bottom line because my markdown estimates are inordinately large; rather, it's because labor itself is such a large expense for firms in the first place, such that modest markdowns still have a meaningful impact on a firm's bottom line.

IA.10 Discussion of Alternative Explanations

IA.10.1 Reconciling empirical findings with Hartman-Glaser et al. (2019) and Kehrig and Vincent (2021)

Papers by Hartman-Glaser et al. (2019) and Kehrig and Vincent (2021) discuss related patterns in labor shares to what I find in this paper, but argue for different economic mechanisms. In this section I briefly discuss why the Hartman-Glaser et al. (2019) mechanism couldn't explain the results within my paper's sample period, and also why my findings are not out of line with Kehrig and Vincent (2021), who argue that their decomposition of aggregate labor share changes are more consistent with product market power.

Hartman-Glaser et al. (2019) also focus on Compustat firms, and they similarly show that the share of output accruing to capital owners has increased by the most for highly productive firms. They also show that firm-level idiosyncratic risk rose substantially over the 1960-2014 period and propose a model based on optimal contracting to explain the capital share dynamics. In their model skilled workers demand wage contracts with embedded insurance for bearing firm-specific risk, and productive firms are able to allocate more output to capital. Increasing idiosyncratic risk amplifies this mechanism, leading to a drop in the aggregate labor share driven by productive firms in the right tail of the productivity distribution.

They calibrate the model to match changes in the labor share between the 1960-1970 and 1990-2014 periods. In appendix Figure IA.7 I plot time trends in their idiosyncratic risk measures before and during my sample period. Although firm-specific risk has indeed risen over that time horizon, idiosyncratic risk has actually been flat or slightly declining over the 1991-2014 period that my data covers. Therefore while their mechanism can speak to broad patterns in labor shares over a longer time horizon, it is incomplete as an explanation for the change in the aggregate labor share *within* the 1991-2014 period that is the focus of this

paper.


[Kehrig and Vincent \(2021\)](#) document that large and productive manufacturing firms with low labor shares have driven the aggregate decline in the manufacturing sector labor share. They argue that this is more consistent with demand side forces than with wage markdowns, because in their framework monopsony should primarily depress labor shares by reducing wages. However, I find that monopsony power is larger for productive firms precisely because they can pay high wages, which allows them to operate on a relatively inelastic portion of their labor supply curve. From the perspective of my findings, the demand side forces they find exaggerate productivity advantages and increase the spread in labor market power. Thus the results in this paper and those in [Kehrig and Vincent \(2021\)](#) are compatible explanations for the aggregate labor share decline.

IA.11 Appendix Figures and Tables


Appendix Figures

Figure IA.1: Compustat-LEHD Matched Sample: Shares of Employment, Market Cap and Sales by Year

Panel A: Matched Sample Shares for Entire Compustat Universe




Panel B: Matched Sample Shares for Non-Financial and Non-Utilities Industries




Note: This figure gives the shares of total Compustat sales, employment, and stock market cap that are represented in my Compustat-LEHD matched sample. The sample period spans 1991-2014. In panel A compares the matched sample shares to the entire Compustat universe; panel B compares the matched sample shares relative to all Compustat firms in non-financial (2-digit NAICS code 20) and non-utilities (2-digit NAICS code 22) industries, since these industries are dropped from the sample.

Figure IA.2: Firm-Level Capital Growth Response to a Firm-Specific Stock Return Shock




Note: This figures shows the responses of firm total capital (physical plus intangible) to a stock return shock, as in (10) in the main text for $h = -5$ to 5 years. See main text for further details.

Figure IA.3: Dynamics of employment, wages, capital, and TFP: calibrated dynamic monopsony model vs data




Note: This figure compares model-implied growth rates in average labor, wages, capital, and employment over the given horizon and using the calibration of the dynamic model described appendix section [IA.1.1](#). Model-generated outcomes are calculated in the exact same way as their empirical counterparts. The analogous empirical estimates along with 95% confidence intervals are in red. Responses are scaled relative to the initial productivity shock, so that the TFP response in the first period is equal to 1. See [IA.1.1](#) for further details.

Figure IA.4: Aggregate Labor Shares Over Time




Note: This figure shows the aggregate labor shares for firms in my Compustat-LEHD matched sample. The dashed lines give the mean aggregate labor share for the given measure over the 1991-2002 and 2003-2014 subperiods. I report aggregate labor shares using two measures of firm labor expenses. Labor Share (LABEX) uses my LEHD-based measure of firm labor expenses, while Labor Share (Predicted XLR) is constructed by imputing Compustat staff and labor expenses. I describe my imputation of XLR in section A.3 of the appendix. Labor shares are computed as the ratio of total labor expenses to total value-added, where I define value-added following Donangelo et al. (2019).

Figure IA.5: Spread in Average Log VA/Worker, and Log Tobin's Q (Total) Between High- and Low-Labor Productivity Firms Trends Upward Over Time




Note: This figure shows the differences between firms in the top- and bottom-quartiles of productivity for the average log value-added per worker; firm worker skill level is from (A.8) in the text and log total Tobin's Q ratio is from Peters and Taylor (2017). Sample period spans 1991-2014.

Figure IA.6: Elasticity Estimates By Labor Productivity for Firms With High- and Low Reliance on SG&A Expenses (Proxy for Stock-based Compensation Intensity)



Note: This figure shows supply elasticity estimates from estimating a version of (13), where firms are double-sorted on the ratio of sales, general, and administrative expenses to LEHD labor expenses and labor productivity. Following Eisefeldt et al. (2021), this measure is used as a proxy for intensity of reliance on equity-based compensation and other perks for skilled employees. I first group firms into above- and below-median bins on the ratio of SG&A to LEHD expenses, and then group them into quartiles on log value-added per worker within these bins. Since accounting practices for reporting SG&A vary by industry (Eisefeldt and Papanikolaou, 2013), I compute the SG&A rankings within 2-digit NAICS category. Elasticities for high SG&A to labor exp firms are in blue; low ranking firms are in red. I include 95% confidence intervals for these estimates based on standard errors double-clustered by industry and year; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages.

Figure IA.7: Idiosyncratic Risk Has Increased Over the Long Term, But is Declining or Flat Within the 1991-2014 Period



Note: This figure plots the log average idiosyncratic stock return and sales growth volatility from [Hartman-Glaser et al. \(2019\)](#). The data span 1950-2014 for returns and 1965-2014 for sales. The 1991-2014 period covered in my paper is shown in red and the period before is shown in blue. Dashed lines give a separate linear time trend for the 1991-2014 and pre-1991 periods.

Appendix Tables

Table IA.1: Compustat Matched Sample and Overall Compustat Summary Stats

Panel A: LEHD-Compustat Matched Sample						
	N	Mean	SD	P5	P50	P95
Log Assets	64000	5.862	1.946	2.887	5.752	9.327
Log Market Cap	63000	12.690	2.070	9.457	12.610	16.280
Log Sales	63500	0.646	1.861	-2.254	0.588	3.845
Log Employment	64000	5.872	2.009	2.769	5.843	9.260
Log Phys Capital	64000	4.914	2.205	1.542	4.808	8.734
Log Int Capital	63500	5.058	1.981	2.099	4.924	8.501
Excess Return	61500	0.119	0.542	-0.717	0.108	0.982
Panel B: Full Compustat Sample						
	N	Mean	SD	P5	P50	P95
Log Assets	110000	5.396	2.202	1.991	5.246	9.334
Log Market Cap	164000	12.170	2.076	8.930	12.060	15.790
Log Sales	133000	-0.136	2.231	-3.689	-0.211	3.648
Log Employment	142000	5.180	2.408	1.380	5.133	9.200
Log Phys Capital	122000	4.428	2.643	0.366	4.284	8.992
Log Int Capital	137000	4.318	2.287	0.858	4.172	8.312
Excess Return	159000	0.104	0.518	-0.700	0.092	0.924

Note: This table provides basic summary stats for the full Compustat sample and my LEHD-Compustat merged sample. The sample period spans 1991-2014. Observation counts are rounded in accordance with Census disclosure requirements.

Table IA.2: Elasticity estimate robustness checks: controlling for markups; holding the set of LEHD states fixed; and adjusting for hourly earnings

Elasticity Estimate:	By Time Period		By Productivity	
	1991–2002	2003–2014	Low Pro- ductivity	High Pro- ductivity
Control for Markups	2.913 (0.243)	1.845 (0.148)	3.528 (0.568)	1.688 (0.184)
Baseline, Markups Sample	2.922 (0.298)	1.883 (0.173)	3.528 (0.566)	1.667 (0.173)
Fix LEHD States	2.468 (0.252)	1.754 (0.221)	2.949 (0.562)	1.691 (0.250)
Baseline, Fix States Sample	2.527 (0.213)	1.810 (0.181)	3.333 (0.635)	1.660 (0.198)
Hours Adjustment			2.860 (0.508)	1.250 (0.212)
Baseline, Hours Adjustment Sample			2.821 (0.440)	1.239 (0.198)

Note: This table reports supply elasticity estimates under various alternative specifications. The “Control for Markups” estimates control for contemporaneous and lagged markup growth (as in [De Loecker et al. \(2020\)](#)). The “Fix LEHD States” specifications back out firm-level wages from the LEHD in the exact same manner as in other tables, except only including workers employed in the set of states found in the LEHD at the end of 1991. The “Hours Adjustment” specification uses workers’ reported usual weekly hours in the American Community Survey (ACS) to estimate supply elasticities in terms of total employment hours and hourly wages. The hours adjustment requires a firm to match to least 20 workers in the ACS in a given year in order to be in the sample that year, and ACS hours are available only from 2005 and on. Elasticity estimates are either heterogeneous by first- and second-halves of the sample period or by above- or below-median productivity in the cross-sectional distribution. Standard errors in parentheses are clustered by industry and year in full sample specifications and industry–year in the subperiod specifications. The firm-year observation counts are as follows: $N = 43,500$ for markups sample; 28,000 for the fix LEHD states sample; and 11,000 for the hours adjustment sample (all rounded in accordance with Census disclosure guidelines).

Table IA.3: Labor Supply Elasticity Estimates Are Insensitive to Additional Controls for Common Market Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Homogeneous									
Elasticity	2.525 (0.305)	2.397 (0.291)	2.492 (0.300)	2.536 (0.289)	2.566 (0.289)	2.568 (0.307)	2.585 (0.303)	2.59 (0.312)	2.614 (0.317)
Panel B: Productivity Sorts									
Elasticity (Low Productivity)	3.525 (0.574)	3.459 (0.560)	3.516 (0.569)	3.572 (0.549)	3.560 (0.547)	3.551 (0.572)	3.541 (0.562)	3.566 (0.575)	3.564 (0.576)
Elasticity (High Productivity)	1.683 (0.173)	1.600 (0.166)	1.677 (0.169)	1.674 (0.161)	1.688 (0.166)	1.706 (0.172)	1.722 (0.171)	1.724 (0.174)	1.727 (0.182)
P-value (High - Low)	0.001	0.000	0.001	0.000	0.000	0.001	0.001	0.001	0.001
Base Controls	X	X	X	X	X	X	X	X	X
Competitor Emp Growth			X					X	X
Competitor Wage Growth			X					X	X
Local Controls + Comp. Return									X
Year FE		X	X						
Industry $\times$ Year FE	X					X	X	X	X
Labor Market $\times$ Year FE				X	X*	X	X*	X	X

Note: This table shows estimates of (9) in the main text when different sets of controls are considered. Panel A considers homogeneous supply elasticities and panel B allows for supply elasticities to vary based off above- and below-median labor productivity status. Baseline control variables include the contemporaneous change in AKM worker effects (worker skill), and lagged wage, employment, and asset growth at the firm level, and my estimate for the baseline specification is in the first column. Industry fixed effects are defined at the 3-digit NAICS level. Labor market fixed effects are from empirically defined labor market clusters (estimation described in section IA.8 of the appendix) with $K = 10$ labor market clusters per year. The “X*” in the 5th and 7th columns indicates that I alternatively use $K = 20$ labor market clusters per year. Competitor emp and wage growth are an average of the growth rates of competitor firms that either hire from or whose employees are hired by the given firm. Comp. return is the average excess stock return of the same competitors. Local market controls include the employment-weighted average changes in local labor market concentration and unemployment rates in the commuting zones across all the firm’s establishments, as well as the average stock returns of other public firms who are headquartered in the same commuting zone. Wage data are from the LEHD, and the sample period spans 1991-2014. See appendix section IA.6.1 for more details. Standard errors double clustered by industry and year are in parentheses; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages.

Table IA.4: Labor supply elasticity estimates: local labor market sorts

Sorting variable rank:	Quartile 1	Quartile 2	Quartile 3	Quartile 4	P-val 4-1
Sort on Log VA/Worker Relative to Local Average					
Elasticity	2.31 (0.26)	1.59 (0.21)	1.41 (0.19)	1.06 (0.14)	0.000
Sort on Local Wage Bill Share					
Elasticity	1.80 (0.21)	1.63 (0.17)	1.64 (0.16)	1.27 (0.19)	0.044

Note: This table reports IV estimates of labor supply elasticities obtained from estimating appendix equation (IA.71). The specification allows for analyzing employment and wage responses at a more granular level than a Compustat gvkey. The unit of analysis is an LEHD SEIN-year, and employment and average wages are taken at the SEIN level. An SEIN in the LEHD corresponds to a state tax reporting unit, and is typically somewhere between a firm and establishment. Each SEIN is assigned an industry and county in the LEHD (these are the modal industry and county if the SEIN has multiple establishments). See appendix section IA.6.2 for more information on LEHD SEINs. SEINs are sorted into quartiles each year either based on their labor productivity advantage compared to firms in the same local labor market (three-digit NAICS interacted with commuting zone) in the top panel, and on their wage bill share of the local labor market in the bottom panel. Specifications now include local labor market times year fixed effects, in addition to the standard set of controls used in Table 3 in the main text. The average local wage bill share of firms the top quartile in the bottom panel is 12.45%. The number of observations is 421,000 (rounded for Census disclosure purposes). Standard errors clustered by Compustat gvkey and year are in parentheses.

Table IA.5: Association between elasticity-implied log markdowns and firm outcomes: local market sorts

Panel A: Markdowns Implied by Elasticity Estimates from Sorting on Log VA/Worker Relative to Local Average								
	Log Lshare		Q (Tot)		ROA_t		ROA_{t+1}	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Markdown	1.06 (0.116)	1.01 (0.108)	-3.00 (0.21)	-1.73 (0.18)	-0.41 (0.033)	-0.35 (0.038)	-0.32 (0.025)	-0.165 (0.016)
Log Markup	-0.225 (0.062)	-0.357 (0.108)	0.87 (0.15)	0.61 (0.15)	0.062 (0.013)	0.095 (0.026)	0.063 (0.012)	0.058 (0.015)
N	54000	53500	51000	50000	54000	53500	50500	49500
R^2 (within)	0.33	0.32	0.19	0.17	0.16	0.12	0.13	0.063
Panel B: Markdowns Implied by Elasticity Estimates form Sorting on Local Wage Bill Share								
	Log Lshare		Q (Tot)		ROA_t		ROA_{t+1}	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Markdown	0.461 (0.124)	-0.127 (0.087)	-0.325 (0.346)	-1.39 (0.348)	-0.118 (0.030)	-0.074 (0.031)	-0.146 (0.023)	-0.045 (0.034)
Log Markup	-0.351 (0.074)	-0.458 (0.119)	0.534 (0.106)	0.282 (0.109)	0.096 (0.015)	0.083 (0.019)	0.081 (0.013)	0.039 (0.012)
N	52500	51500	53500	52500	57500	57000	53500	53000
R^2 (within)	0.27	0.25	0.14	0.16	0.25	0.10	0.22	0.06
Size Controls	X	X	X	X	X	X	X	X
Capital Intensity Controls	X	X	X	X	X	X	X	X
Ind $\times$ Year FE	X	X	X	X	X	X	X	X
Firm FE Firm FE		X		X		X		X

Note: This table reports results from regressing firm log labor shares, valuations, and contemporaneous and future return on assets on elasticity-implied log wage markdowns and log markups. Implied firm wage markdowns are computed by taking the wage bill weighted-average elasticity of a Compustat gvkey across all SEINs within a gvkey, using elasticity estimates from the appendix Table IA.4. The top and bottom panels respectively take the supply elasticity estimates from the top and bottom panels, respectively, of appendix Table IA.4. Capital intensity controls include the logs of physical and intangible capital per worker, while size controls include the logs of assets, sales, and employment. Markups are computed as in De Loecker et al. (2020). Standard errors clustered by industry and year in parentheses.

Table IA.6: Mean (median) real marginal adjustment costs and wages, by firm productivity and worker skill

	Real Adjustment Cost Per Hire	Average Real Wage	$\frac{\text{Adj. Cost}}{\text{Wage}}$
Panel A: Homogeneous Worker Types			
Overall	\$39,780 (\$28,490)	\$68,470 (\$60,500)	0.58 (0.47)
Overall, 1991-2002	\$27,990 (\$19,590)	\$61,470 (\$55,450)	0.46 (0.35)
Overall, 2003-2014	\$53,810 (\$40,870)	\$77,390 (\$68,450)	0.70 (0.60)
Low Productivity	\$27,610 (\$20,570)	\$49,530 (\$46,030)	0.56 (0.45)
High Productivity	\$51,960 (\$38,190)	\$84,540 (\$74,290)	0.61 (0.51)
Panel B: Skill Heterogeneity			
Low Skill	\$788 (\$509)	\$29,770 (\$30,170)	0.03 (0.02)
Middle Skill	\$4,137 (\$2,804)	\$54,610 (\$54,130)	0.08 (0.05)
High Skill	\$218,200 (\$142,000)	\$142,300 (\$131,700)	1.53 (1.08)

Note: This table reports estimates of mean and median marginal adjustment costs (in real 2011 US dollars) across all firm-years for the given set of firms or subperiods, and following the procedure outlined in appendix sections [IA.1.3](#) and [IA.1.4](#). Medians are reported in parentheses. The second column reports that average real wage (also in real 2011 US dollars) for the same set of firms, while the third column divides the mean (median) marginal adjustment cost by the mean (median) real wage. Marginal adjustments costs are calculated as the product of estimated adjustment cost parameter γ_L and the firm's gross hiring rate for the year (defined as total new hires in the current year divided by total workers in the previous year). Hiring rates are winsorized at the 1% level. Panel A imposes a single γ_L for all firm-years within a given productivity ranking and time period (1991-2002 versus 2003-2014). Panel B allows for γ_L to also vary by worker skill as defined in [Tables 1](#) and [3](#) and the main text.

Table IA.7: Labor supply elasticity estimates for quantification: by period, productivity, and skill

	Low Productivity				High Productivity			
	Overall	Low Skill	Middle Skill	High Skill	Overall	Low Skill	Middle Skill	High Skill
Elasticity (1991–2002)	4.05 (0.64)	11.25 (2.87)	7.77 (1.40)	1.95 (0.26)	1.94 (0.14)	6.36 (0.96)	4.97 (0.53)	1.28 (0.11)
Elasticity (2003–2014)	2.67 (0.29)	5.77 (0.90)	5.19 (0.73)	1.19 (0.14)	1.21 (0.12)	4.37 (0.48)	3.58 (0.46)	0.72 (0.10)

Note: This table reports the IV labor supply elasticity estimates by time period, labor productivity rank, and worker skill that are used to generate estimates of Π_t in (19) in the main text. Standard errors in parentheses are clustered by industry–year, and controls are otherwise the same as defined in Table 3 of the main text.

Table IA.8: Productive Firms Face Lower Supply Elasticities For Workers of All Skill Levels—Within Industry Productivity Sorts

Productivity:	Quartile 1	Quartile 2	Quartile 3	Quartile 4	P-val 4-1	R-sq	N
Whole Firm							
Employment	0.11 (0.01)	0.10 (0.01)	0.11 (0.01)	0.11 (0.01)	0.641	0.13	43500
Wages	0.03 (0.004)	0.04 (0.003)	0.05 (0.004)	0.09 (0.01)	0.000	0.42	43500
Elasticity	4.03 (0.79)	2.62 (0.38)	2.24 (0.28)	1.23 (0.19)	0.001		
Low Skill Workers							
Employment	0.12 (0.01)	0.12 (0.01)	0.13 (0.02)	0.14 (0.01)	0.285	0.11	43500
Wages	0.01 (0.003)	0.02 (0.003)	0.02 (0.002)	0.03 (0.004)	0.016	0.29	43500
Elasticity	8.79 (2.12)	6.84 (1.69)	6.66 (1.06)	4.95 (0.64)	0.122		
Middle Skill Workers							
Employment	0.11 (0.01)	0.10 (0.01)	0.11 (0.01)	0.10 (0.01)	0.444	0.08	43500
Wages	0.02 (0.002)	0.02 (0.002)	0.02 (0.002)	0.03 (0.003)	0.005	0.18	43500
Elasticity	7.28 (1.43)	5.28 (0.87)	5.74 (0.88)	3.58 (0.46)	0.031		
High Skill Workers							
Employment	0.10 (0.01)	0.09 (0.01)	0.09 (0.01)	0.09 (0.01)	0.266	0.08	43500
Wages	0.05 (0.01)	0.06 (0.004)	0.08 (0.01)	0.12 (0.01)	0.000	0.45	43500
Elasticity	2.05 (0.35)	1.46 (0.20)	1.25 (0.14)	0.74 (0.13)	0.002		

Note: This table contains supply elasticity estimates for firms sorted on log value-added/worker quartiles as in Table 3, except I now sort firms on labor productivity within 2-digit NAICS industry. Controls include 3-digit NAICS industry by year and productivity quartile fixed effects; lagged growth rates in wages, employment, and total assets; and the contemporaneous change in average worker skill level at the firm (see (A.8) for definition). Workers are placed into skill groups based on their estimated worker effects from a modified Abowd et al. (1999) style wage decomposition with time-varying firm fixed effects. Individuals in the bottom two quintiles of the cross-sectional distribution of worker effects are considered low-skilled, the third and fourth quintiles middle-skilled, and the top quintile high-skilled. Changes in average worker skill are computed within the population of workers considered in the specification. “P-val 4-1” gives the p-value from a test that the coefficients for firms in the top and bottom quartiles have equal values. Wage data are from the LEHD, and the sample period spans 1991-2014. Standard errors clustered by industry and year are in parentheses; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages. See section 3 in main text for more details.

Table IA.9: Elasticities for Firms Sorted on Productivity—Robustness

Productivity:	Quartile 1	Quartile 2	Quartile 3	Quartile 4	P-val 4-1	R-sq	N
Panel A: 3-Year Horizon							
Employment	0.20 (0.02)	0.17 (0.02)	0.20 (0.02)	0.19 (0.02)	0.563	0.20	34000
Wages	0.03 (0.002)	0.04 (0.002)	0.05 (0.004)	0.08 (0.01)	0.000	0.51	34000
Elasticity	7.78 (0.88)	4.79 (0.54)	4.42 (0.63)	2.44 (0.34)	0.000		
Panel B: Wage Growth Using Changes in AKM Firm Effects							
Employment	0.12 (0.01)	0.09 (0.01)	0.12 (0.01)	0.11 (0.01)	0.240	0.13	43500
Wages	0.02 (0.002)	0.02 (0.002)	0.03 (0.002)	0.04 (0.005)	0.000	0.25	43500
Elasticity	6.95 (1.15)	4.24 (0.48)	4.07 (0.45)	2.47 (0.37)	0.003		
Panel C: TFP Sort							
Employment	0.11 (0.01)	0.09 (0.01)	0.10 (0.01)	0.12 (0.01)	0.482	0.14	37500
Wages	0.03 (0.004)	0.04 (0.004)	0.05 (0.01)	0.08 (0.01)	0.000	0.43	37500
Elasticity	3.21 (0.64)	2.19 (0.33)	1.91 (0.26)	1.42 (0.21)	0.007		
Panel D: LBD Employment Growth							
Employment	0.09 (0.01)	0.08 (0.01)	0.10 (0.01)	0.09 (0.01)	0.683	0.08	40000
Wages	0.03 (0.003)	0.03 (0.003)	0.05 (0.004)	0.09 (0.01)	0.000	0.42	43500
Elasticity	3.10 (0.75)	2.55 (0.33)	1.85 (0.28)	1.01 (0.17)	0.01		

Note: This table shows robustness checks for the elasticity estimates obtained from estimating variants of (13) in the main text. All equations use the baseline set of controls from (13), which includes industry $\times$ year fixed effects. Panel A estimates a variant of (13) where stock returns and employment/wage growth are at the 3-year horizon. In Panel B changes in AKM firm effects $\phi_{j,t}$ from estimating (A.5) replace growth in the firm-level average wage in estimating the supply elasticity. Panel C sorts on firm total-factor productivity from İmrohoroglu and Tüzel (2014) instead of log value-added per worker. Panel D uses employment growth from Longitudinal Business Database employment figures instead of Compustat. “P-val 4-1” gives the p-value from a test that the coefficients for firms in the top and bottom quartiles have equal values. Wage data are from the LEHD, and the sample period spans 1991-2014. Standard errors clustered by industry and year are in parentheses; elasticity standard errors are estimated via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages. See section 3 in main text for more details.

Table IA.10: Asymmetric Wage Passthrough for Positive and Negative Stock Return Shocks

Asymmetric Wage Responses (All Workers)				
Productivity:	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Wages (Ex Ret ≥ 0)	0.031 (0.006)	0.037 (0.006)	0.056 (0.007)	0.092 (0.010)
Wages (Ex Ret < 0)	0.021 (0.005)	0.024 (0.007)	0.040 (0.005)	0.089 (0.016)

Note: This table shows estimates for asymmetric wage responses to positive and negative excess stock return shocks for firms sorted on labor productivity, as in equation (IA.73). See appendix section IA.7.2 for further details.

Table IA.11: Labor Supply Shocks from Union Elections and Changes in Non-Compete Enforceability Do Not Affect Supply Elasticity Estimates

Panel A: Union Elections			
	(1)	(2)	(3)
Elasticity	1.65 (0.33)	1.64 (0.32)	1.72 (0.34)
P-Val (Diff)		0.57	0.35
N	4200	4200	4200
Baseline Controls	X	X	X
Union Dummies		X	X
Union Dummies $\times$ Excess Return			X
Industry $\times$ Year FE	X	X	X
Union Win Excess Return	-0.08***		
Panel B: Non-Compete Changes			
	(1)	(2)	(3)
Elasticity	2.37 (0.50)	2.36 (0.50)	2.4 (0.52)
P-Val (Diff)		0.29	0.70
N	11000	11000	11000
Baseline Controls	X	X	X
Non-Compete Dummies		X	X
Non-Compete Dummies $\times$ Excess Return			X
Industry $\times$ Year FE	X	X	X

Note: Panel A of this table shows how elasticity estimates change when including dummies for a union election occurring and for a union win in years t , $t - 1$, or $t + 1$. Elections are taken from a list compiled by [Knepper \(2020\)](#). The third column adds interactions of union election dummies with the excess return in that year. Panel B of this table uses non-compete changes from the lists compiled by [Jeffers \(2019\)](#) and [Ewens and Marx \(2017\)](#). Following [Ewens and Marx \(2017\)](#) firms are considered treated if their headquarters state changes non-compete laws in a given year; non-compete controls include separate dummies for non-compete increases and decreases in the years t , $t - 1$, and $t + 1$, and the third column interacts these dummies with stock returns. “P-Value (Diff)” gives the p-value on the differences between the elasticity using the given shock and for my baseline estimate. The sample is restricted to only those firms who ever had a unionization or non-compete event, and the baseline estimate is computed for the same sample. Standard errors double clustered by industry and year are in parentheses, and are computed by estimating the elasticity via two-stage least squares where wages are predicted in the first stage and employment is then regressed on predicted wages.

Table IA.12: Productive Firms Hire More Skilled Workers on Average

Dep Var: Log Firm Average AKM Worker Effects				
log VA/Worker	0.189 (0.011)	0.147 (0.018)	0.163 (0.011)	0.110 (0.014)
Size Controls		X		X
Industry X Year FE			X	X
N	57500	57500	57500	57500
R ² (within)	0.344	0.386	0.217	0.263

Note: This table shows regressions of log firm average AKM worker fixed effects (defined in (A.8)) on firm labor productivity as proxied by log value-added per worker. Controls for size include log employment, assets, and sales. Standard errors double clustered by year and industry in parentheses.

Table IA.13: Explanatory Power of Labor Market Proxies for Worker Flows, Wages, Stock Returns, and Employment

	K = 10	K = 20	NAICS2	NAICS3	K = 10 + NAICS3
Worker Flow Share	0.50	0.42	0.34	0.23	
Log Wage	0.62	0.65	0.51	0.60	0.70
Wage Growth	0.03	0.04	0.03	0.03	0.04
Excess Return	0.14	0.15	0.11	0.16	0.18
Log Emp	0.19	0.26	0.10	0.18	0.26
Emp Growth	0.03	0.03	0.03	0.03	0.04

Note: This table shows the explanatory power of different candidate labor market boundaries for several different variables. The first row reports the fraction of worker transitions between Compustat firms that occur within the same candidate market definition. The remaining rows report the adjusted R^2 from a regression of the given variable on market $\times$ year fixed effects. The first two columns show results for labor market clusters estimated with 10 or 20 clusters and the next two columns instead use either 2- or 3-digit NAICS codes. The last column reports adjusted R^2 values for labor market $\times$ year FEs and 3-digit NAICS $\times$ year FEs included simultaneously. Sample spans 1992-2013.

Table IA.14: Productive Firms Pay Higher Wages and have Lower Separations Rates for Workers of all Skill Levels

Productivity:	Quartile 2	Quartile 3	Quartile 4
Panel A: Log Wages			
Whole firm	0.186 (0.028)	0.366 (0.028)	0.613 (0.034)
Low skill	0.139 (0.017)	0.221 (0.017)	0.298 (0.014)
Middle skill	0.098 (0.013)	0.174 (0.012)	0.265 (0.011)
High skill	0.099 (0.015)	0.188 (0.017)	0.349 (0.025)
Panel B: Separations Rates			
Whole firm	-0.075 (0.004)	-0.104 (0.005)	-0.125 (0.006)
Low skill	-0.069 (0.004)	-0.093 (0.005)	-0.10 (0.006)
Middle skill	-0.066 (0.004)	-0.089 (0.005)	-0.105 (0.006)
High skill	-0.055 (0.004)	-0.073 (0.006)	-0.088 (0.007)

Note: This table shows coefficients on labor productivity quartile dummies α_q from regressions of the form

$$y_{j,t} = \alpha_{q(j,t)} + \alpha_{I(j),t} + \epsilon_{j,t} \quad (\text{IA.76})$$

for firm j productivity quartile $q(j,t)$ at time t . The outcome variable $y_{j,t}$ is the average log wage or separations rate for the given worker type. $\alpha_{I(j),t}$ denote industry-year fixed effects. The bottom productivity quartile bin is the omitted category. Workers are grouped into skill groups based on their estimated worker effects from a modified [Abowd et al. \(1999\)](#) style wage decomposition with time-varying firm fixed effects. Individuals in the bottom two quintiles of the cross-sectional distribution of worker effects are considered low-skilled, the third and fourth quintiles middle-skilled, and the top quintile high-skilled. Standard errors double clustered by year and industry in parentheses.