Valuing Labor Market Power: The Role of Productivity Advantages

Bryan Seegmiller*

First Version: November 2021

Current Version: December 2023

Abstract

Public firms with high labor productivity have a large and expanding labor market competitive advantage. Using firm-specific stock returns to estimate heterogeneous labor supply elasticities by labor productivity and across time, calibrated to a dynamic wage posting model featuring costly hiring, I estimate wage markdowns which largely explain: a wide cross-sectional labor share spread by productivity; the public firm aggregate labor share decline from 1991-2014; and productive firms' high valuations given their modest investment rates. Cashflows from wage markdowns are worth two-fifths of aggregate capital income. Market power over skilled workers may play an important role in these patterns.

*Kellogg School of Management. I thank my advisors Leonid Kogan, David Autor, Larry Schmidt, and Dimitris Papanikoloau for extremely helpful guidance and continued support. I also thank Daron Acemoglu, Hui Chen, Taha Choukhmane, Maryam Farboodi, Dan Greenwald, Simon Jager, 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, and CBDRB-FY22-SEHSD003-004).

Market power affects the distribution of firm cashflows between the various claimants on its output: increased market power shifts the distribution of the value of productive output towards owners of the firm's capital, at the expense of the firm's workforce. Empirical patterns that have emerged over the past several decades suggest that US firms have been earning higher rents due to increased market power. This period of time has been characterized by a large rise in the value of the aggregate stock market and high firm profitability (Greenwald, Lettau, and Ludvigson, 2021; Barkai, 2020); high firm valuations but with low rates of investment (Gutierrez and Philippon, 2017); and, a decline in the share of output accruing to labor (Karabarbounis and Neiman, 2013).

Much of the literature documenting the role of market power in explaining these trends has focused on product markets (De Loecker, Eeckhout, and Unger, 2020; Autor, Dorn, Katz, Patterson, and Van Reenen, 2020). In this paper I show that firms with high labor productivity enjoy a competitive advantage in the labor market in the form of large wage markdowns from marginal revenue product, and I quantify the value of the cashflows derived from this advantage. This "superstar firms" view of labor market power can explain cross-sectional differences in firm-level profitability, labor shares, and firm valuation ratios, as well as time series changes in the aggregate labor share, in a sample of publicly-traded US corporations which contains both wide variation in labor productivity and features the largest and most highly productive firms in the world. Labor compensation is the single largest cost of production for most firms, which means that paying workers below their marginal revenue product can generate considerable value to capital owners: holding equilibrium quantities fixed, I find that the value of wage markdowns represents roughly a third of the capital income of the typical US corporation, and even more for the most productive firms.

My estimates of labor market power also track important macroeconomic trends in the competitive landscape of the US economy. An expanding labor productivity gap between the most- and least-productive firms parallels an absolute and relative increase in productive firms' labor market power—as measured by the supply elasticity facing the firm—contributing to a larger cross-sectional spread in labor shares and a lower aggregate labor share. In contrast with the estimates of labor market power directly implied by my supply elasticity estimates, measures of local labor market concentration have trended downward, and hence cannot explain these aggregate trends.[1] Additionally, in the cross-section, productive firms have much higher valuation ratios but they do not invest more; much of this cross-sectional "investment gap" can be explained by accounting for the capital income productive firms derive from wage markdowns.

I build towards these conclusions as follows. I first make a unique contribution to the burgeoning literature on the passthrough of firm-specific shocks to labor outcomes by studying

[1] Berger, Herkenhoff, and Mongey (2021), Rinz (2018), and Lipsius (2018) all find declining local labor market concentration over the past several decades using Census administrative data.

the heterogeneous impact of idiosyncratic stock returns on firm-level labor demand by worker and firm type using employer-employee matched data from the Longitudinal Employer-Household Dynamics (LEHD) matched to publicly traded firms in Compustat. Despite a massive body of prior research studying patterns associated with stock returns, my paper is the first to ever offer such a systematic analysis of the relationship between a firm's own stock returns and its labor demand.[2] Consistent with stock returns containing unexpected information about labor demand, I show in a series of local projection specifications that firm-specific stock returns induce immediate, persistent, and highly positive effects on the forward growth rates in firms' average wages and employment levels, while they do not predict past growth rates. Stock returns also correlate with immediate improvements in value-added per worker, which is directly proportional to Cobb-Douglas marginal revenue productivity, bolstering the labor demand shock interpretation.

These patterns are broadly consistent with models of imperfect competition in frictional labor markets where firm-specific shocks are impounded into wages (see Card, Cardoso, Heining, and Kline (2018) for a broad overview of this literature); I further examine dimensions of heterogeneity which can better highlight specific models which are consistent with the data. In wage posting models the ratio of the employment and wage responses to a labor demand shock can be interpreted as a firm-specific labor supply elasticity. 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, and lower elasticities for incumbent workers compared to new recruits (1.7 versus 7.2) and for high labor productivity firms (4.3 for firms in the bottom labor productivity quartile and 1.2 for firms in the top quartile). The ratio of the wage growth to value-added growth responses to stock returns yields a rent-sharing elasticity; I obtain a pooled average rent-sharing elasticity of 0.25, similar to prior work using firm-level aggregates (Abowd and Lemieux, 1993; Van Reenen, 1996) and also close to the estimate of 0.35 found in Kline, Petkova, Williams, and Zidar (2019). Examining heterogeneity in rent-sharing elasticities, they are higher for incumbents than recruits (0.29 versus 0.15) and increasing by firm labor productivity rank (0.16 in the bottom quartile and 0.35 in the top quartile).

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 power. Moreover, stylized facts on firms with higher labor productivity also 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. My estimates allow me to quantify the role of labor market power in explaining these

[2] While Kogan, Papanikolaou, Schmidt, and Song (2020) correlate individuals' wage growth with stock returns, it is not their primary focus, which is instead on technological innovation and income risk.

differences. Second, many models of imperfect competition in labor markets directly imply increasing rent-sharing elasticities as value-added per worker rises, which implication has not been closely tested in prior work. My estimates directly confirm this prediction.

In addition to the heterogeneity by incumbent/recruit status and by labor productivity rank, the strong relationship between stock returns and labor demand also affords me the statistical power to examine how labor supply elasticities have changed over time. I find that across the productivity distribution, supply elasticities are uniformly lower in the second half of my sample period (2003-2014) than the first half of my sample period (1991-2002). Through the lens of a wage posting, monopsonistic model, this implies a secular increase in labor market power, and a role for changes in the economic rents earned from labor market power to explain time-series trends in labor shares in my sample of publicly traded firms. Applying my method at higher frequency using 3-year moving window average estimates, a widening cross-sectional gap in wage markdowns between the most- and least-productive firms implied by time-varying supply elasticity estimates closely tracks the widening gap in the average log labor share, suggesting these estimates capture genuine differences in market power in the labor market.

My reduced-form empirical results are suggestive of the cross-sectional and time-series trends in labor market power, but quantifying their implications for the contribution to firm cashflows and labor shares requires interpretation through a model. I show that the lower estimated labor supply elasticities and higher rent-sharing elasticities as labor productivity goes up are jointly consistent with both a wage posting model where firms face more inelastic supply as their labor demand grows (Card et al., 2018), or an intrafirm bargaining framework with multiworker firms (Stole and Zwiebel, 1996; Acemoglu and Hawkins, 2014) where firms pay a convex cost to hire workers, though the structural interpretation of the supply elasticity estimates differ between the two. Meanwhile, the differential labor demand responses for incumbent workers and recruits point to a model where firms find it costly to replace incumbents and consequently share more rents with them in order to increase retention (Kline et al., 2019). These joint features of the data inform my choice to adopt a dynamic wage posting framework where firms make different wage offers to retain incumbents and attract new recruits. While the wage posting framework is most consistent with the retention motive that appears to be a key aspect of the data, my model delivers predictions that are closely related to the aforementioned bargaining frameworks, and in fact yield predictions that are observationally equivalent under certain assumptions on the convexity of the hiring cost in the bargaining formulation; otherwise, the direction of the predictions on wage markdowns is the same in both cases, and economically reasonable assumptions on the amount of convexity in the hiring cost function in the bargaining framework imply that markdown estimates under a wage posting regime are quantitatively very close or even conservative relative to an alternative bargaining framework.

Motivated by my empirical finding that idiosyncratic stock returns also affect the firm-specific wage of new hires, firms in my model choose wage offers for their recruits. In the model the firm faces a trade off between paying a convex replacement cost required to hire and train new recruits with the higher wage offers needed to increase the probability of retaining incumbent workers, and the number of workers retained until next period becomes a state variable which makes the firm's problem dynamic. These features are both unique to my model relative to Kline et al. (2019). The calibration of my parsimonious model closely mimics key targeted moments in the data, including the pooled average firm-level supply elasticity; the differential average supply elasticities and passthrough of stock returns to wage growth of incumbents and recruits; the average separations rate of incumbent workers; and the empirical wage premium of an incumbent worker over an equivalently skilled recruit. It also captures cross-sectional regularities, delivering quantitatively accurate lower labor supply elasticities, higher rent-sharing elasticities, and lower labor shares for more productive firms.

The model implies that firms reduce their exercise of labor market power considerably relative to a static monopsony model, because applying the static wage markdown rule would imply sub-optimally low retention rates. The model delivers the ratio between the true wage markdown from marginal revenue product and relative to the elasticity-implied markdown in a static model; this key quantity, which I call the dynamic markdown wedge, maps my reduced-form empirical estimates to quantitative wage markdowns.

Combining my empirical results and the model calibration, my main quantitative findings are the following. The full-sample calibration of the dynamic model with hiring adjustment costs shows that wage markdowns are about 16% smaller relative to a static model without costly adjustment. However, the dollar value of wage markdowns still represents a substantial portion of firm cashflows. The average firm pays workers about 83% of their marginal product; firms in the top productivity quartile pay about 62%, while the least productive firms pay 94%. I construct a counterfactual where I remunerate back to workers their implied marginal product of labor, holding firms' production decisions constant, which quantifies the cashflows coming from the gap between wages and marginal product at observed equilibrium quantities. Wage markdowns are worth 34% of the average firm's operating income; this figure is 43% (17%) for firms in the top (bottom) quartile of the labor productivity distribution. In aggregate, wage markdowns are worth about two-fifths of total operating income.

Wage markdowns can also account for about three-fifths of the wide gap in average log labor shares between the firms in the top and bottom quartiles of the labor productivity distribution; my supply elasticity estimates and separate model calibrations for the years 1991-2002 and 2003-2014 further imply that changes in markdowns can explain the majority of the observed decline in the average aggregate labor share among Compustat firms between the two time periods. Finally, productive firms have high valuations without high investment rates. Removing the cashflows coming from wage markdowns explains the majority of cross-sectional differences in average Tobin's Q across firm labor productivity ranks, better aligning productive firms' average valuations with their modest investment rates.

In arriving at all of the above, I interpret my stock return-induced estimates as identifying the true elasticity parameters. I examine thoroughly the conditions under which this interpretation is justified, and what deviations from those conditions mean for my estimates quantitatively. Conditional on market-level controls, my estimates identify the supply elasticity under the assumption that shocks to the level of the labor supply curve only operate at the market level; under the more conservative assumption that firm-specific labor supply shocks do not strongly negatively comove with firm-specific productivity shocks, they provide an upper bound on the elasticity, which causes my estimates to be a conservative estimate of the magnitude of labor market power. The most important potential source of confounding variation comes from common market level productivity shocks, which could bias my supply elasticity estimates downward. Accordingly, I show that, relative to my baseline specification, my elasticity estimates are invariant to the way I control for market shocks, showing remarkable stability under a wide variety of alternative sets of controls for common shocks. This confirms that my estimates are driven by the firm-specific component of stock returns. I also find that alternate labor demand shock proxies—such as using the stock returns of firms' customers, patent grants, or only including stock returns near quarterly earnings announcements—yield supply elasticity estimates that are nearly identical to and statistically indistinguishable from my baseline, both for homogeneous estimates and when sorting on labor productivity, although estimated less precisely relative to my baseline specification. This suggests that bias from confounding firm-specific labor supply shocks does not play a quantitatively important role.

In a final suggestive set of empirical exercises I examine to what extent my estimates vary by worker skill, since this may be informative about underlying labor market frictions: for skilled workers, labor markets are much less local (see Malamud and Wozniak, 2012; Amior, 2020, for example), diminishing the importance of concentration defined over narrow geographic boundaries. Instead, monopsony power over skilled workers may largely operate through firm-specific human capital that narrows the set of substitute employers. On this front, the wage premium for incumbent workers has increased over time, and worker separations rates have decreased, suggesting that it has both become more costly to replace an incumbent worker, and incumbent workers have become less willing to leave their employers. Breaking down my empirical estimates by worker skill, my key qualitative findings in the time-series and the cross-section maintain for workers of all skill levels, but the most skilled workers—who are also employed more intensively at productive firms—appear to play an especially important quantitative role. While my empirical estimates and model don't directly illuminate the exact underlying labor market frictions, this does suggest that studying in detail the labor market frictions skilled workers face could illuminate the micro-foundations behind these aggregate patterns in labor market power. The increased importance of human capital specificity seems a likely contributor.

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.[3] Most related to my paper is Hartman-Glaser, Lustig, and Xioalan (2019), who document that large firms have obtained an increasing share of their sales as capital income over the 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. I argue that this is in large part driven by productive firms marking down wages further from marginal product. I also provide the novel finding that firms generate a substantial portion of their operating income from wage markdowns, which is especially the case for firms with high labor productivity. This competitive advantage generates rents which are reflected in higher financial valuation ratios and but modest investment rates among firms with high labor productivity.

This paper also 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;[4] another related empirical literature has examined how labor market concentration affects wages.[5] 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. Finally, a large literature estimates the passthrough of firm-specific shocks to worker outcomes.[6] 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 curve that the firm faces.

In a related prior paper, Gouin-Bonenfant (2020) 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,

[3]These 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).

[4]See Lamadon, Mogstad, and Setzler (2019), Berger et al. (2021), Kroft, Luo, Mogstad, and Setzler (2020), Bassier, Dube, and Naidu (2020), and Ransom and Sims (2010) for a few examples. Alternatively, Yeh, Macaluso, and Hershbein (2022) and Mertens (2022) 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.

[5]See Rinz (2018), Benmelech, Bergman, and Kim (2018), Schubert, Stansbury, and Taska (2020), and Jarosch, Nimczik, and Sorkin (2019), Azar, Marinescu, and Steinbaum (2022) for example.

[6]See 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).

including a decline in the aggregate labor share. I differ from Gouin-Bonenfant (2020) in providing direct empirical estimates of firm-level monopsony power as implied by labor supply elasticities. My paper is also the first, to my knowledge, to explicitly estimate heterogeneous supply elasticities for firms of different productivity levels both cross-sectionally, across time, and by worker skill level. I also differ by focusing on quantifying the value of the capital income that is derived from wage markdowns. Thus my findings complement and add to those of Gouin-Bonenfant (2020).

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.[7] 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.

Finally, 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.[8] My findings suggest that accounting for imperfect competition in labor markets may be of first-order importance for future research in this area.

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 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 non-qualified stock options. Individuals in the LEHD are linked to their employers through

[7]See 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.

[8]References in asset pricing include Eisfeldt and Papanikoloau (2013), Belo, Lin, and Bazdresch (2014), Donangelo, Gourio, Kehrig, and Palacios (2019), Kuehn, Simutin, and Wang (2017), Liu (2019); see Matsa (2010), Jeffers (2019), Mueller, Ouimet, and Simintzi (2017), Kim (2020), and Shen (2021) for examples in corporate finance.

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.[9] 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.

1.2 Matched Compustat-LEHD Sample

I link firm identifiers in each year of the the LEHD to Compustat records using a crosswalk created by Larry Schmidt.[10] 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, my matched sample covers about 62% of employment, 63% of market cap, and 50% of sales represented in the entirety of Compustat in a given year, which 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, which covers 71% of employment, 79% of market cap, and 62% of sales among the industries include, 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 various measures of firm size such as assets, employment, or market cap. The distributions of annual excess 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.

[9] An overview of LEHD coverage is available at https://www2.census.gov/ces/wp/2018/CES-WP-18-27.pdf.

[10] Thanks to Larry for sharing his crosswalk code.

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 et al., 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, Kramarz, and Margolis (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 Empirical Results

In this section I first discuss my method for estimating supply and rent-sharing elasticities. I then apply my method to estimate heterogeneous elasticities for by worker type and for firms sorted on labor productivity.

2.1 Firm-Specific Shocks and Identification of Labor Supply and Rent-Sharing Elasticities

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 the basic static monopsony framework (Robinson, 1969). The firm has a revenue function $F(L)$ and faces a labor supply function $L(w)$, where w is the wage offer. The firm solves:

$$\max_{w} \qquad F(L(w)) - wL(w) \tag{1}$$

First order conditions give

$$w = \frac{\epsilon}{1+\epsilon} F_L \tag{2}$$

where

$$\epsilon = \frac{dL}{dW}\frac{w}{L} \text{ is the supply elasticity.} \tag{3}$$

Define

$$\frac{\epsilon}{1+\epsilon} \equiv \text{ wage markdown} \tag{4}$$

Thus the wage markdown represents the fraction of labor's marginal revenue productivity that workers obtain in wages. In perfect competition $\epsilon \to \infty$, and the wage equals F_L, the marginal revenue product of labor.

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.2 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), 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 wage posting, monopsonistic setups from Card, Heining, and Kline (2013) and Kline et al. (2019), as well as my wage posting model in section 3; I show in internet appendix subsection IA.2.1 that inferred wage markdowns from applying equation (4) to estimate the wage markdown when firms instead engage in multilateral bargaining are quantitatively close (and likely to be conservative for reasonable parameterizations of convexity in the hiring cost function). In section IA.2.2 of the internet appendix I further show that the rent-sharing elasticities implied by the two models are also closely related to one another.

In what follows I take the wage posting perspective because it is particularly useful for modeling and interpreting the heterogeneous labor demand responses I will document for incumbent workers and recruits, which reflects a retention motive on the part of the firm for which I find considerable empirical support. This is harder to motivate under the bargaining framework, which assumes that worker separation rates are independent of the wages offered.[11] Still, to the extent that such bargaining motives also operate in labor markets, internet appendix IA.2 makes clear that the close relationship between the two frameworks ensures that quantitative estimates for wage markdowns are not likely to be not substantively biased upward by assuming wage posting.

2.2 Stock Returns As a Labor Demand Shock

An ultimate empirical objective of this paper is to quantify cross-sectional and time series variation in wage markdowns for Compustat firms sorted on labor productivity, so my labor demand shock also needs to be available for a wide variety of firms and for a long period of time. As emphasized by Manning (2003), supply elasticities can be identified by shocks to the marginal revenue product of labor, holding the firm-specific labor supply curve constant. In order to estimate supply elasticities, such labor demand shocks should ideally be: 1) a persistent, unanticipated shock to firm productivity; 2) firm-specific; 3) not correlated with

[11] This is the case in (Acemoglu and Hawkins, 2014), Kuehn et al. (2017), and Garin and Silverio (2020), for instance.

shocks to firm-specific labor supply.

In addition to being widely available, firm-specific stock returns also satisfy these criteria for a labor demand shock quite well. In particular, they reflect a revision in expected discounted future cashflows, and hence inherently correlate with changes in current and expected future revenue productivity.[12] Consider the standard asset pricing equation that determines the price of a stock:

$$P_t = E_t \left[\sum_{k=0}^{\infty} M_{t,t+k} D_{t+k} \right]$$

Here D_{t+k} is a cashflow occuring at time $t+k$ and $M_{t,t+k}$ is the time t discount factor for time $t+k$ cashflows. The time t stock return is

$$R_t = \frac{D_t + P_t}{P_{t-1}}$$

Most of the variation in stock returns is firm-specific, and the common components can be removed rather easily. Because stock returns react to new information, they are inherently unexpected, and because they are forward looking they represent persistent changes in the outlook of the firm going forward. All this suggests that stock returns are good candidates for unexpected shocks to firm revenue productivity, which in turn moves firm labor demand.

However, firms may also gain value from the willingness of workers to be employed at the firm for any given wage, which affects the level of the labor supply curve. Unexpected changes in the level of labor supply could bias supply elasticity estimates when using stock returns as an instrument for labor demand shocks. To build intuition about the form this bias may take, I derive closed form expressions for the bias in internet appendix IA.1 in a simple model of the labor market that takes inspiration from Card et al. (2018) and Lamadon et al. (2019). There are two key insights from this exercise. The first is that failing to account for market-wide shocks is likely to bias stock-return based elasticity estimates towards zero, causing me to make erroneously large inferences about the importance of labor market power. Because of this I demonstrate the insensitivity of my baseline elasticity estimates to a host of specifications using different proxies for market-level shocks.

The second insight is that once market shocks are accounted for, any remaining bias

[12] Previous research provides evidence that while discount rate shocks drive much of the variation in the aggregate stock market, news containing information about future cashflows drives the vast majority of the variation in firm-level stock returns. For example, Vuolteenaho (2002) finds that the idiosyncratic component of individual stock returns is almost entirely driven by firm-specific cashflow news, while movements in discount rates are primarily common across firms. More recent research (Neuhierl, Scherbina, and Schiusene, 2013; Boudoukh, Feldman, Kogan, and Richardson, 2018) also shows that fundamental information about firms' cashflows, such as unexpectedly high earnings or new product announcements, are an important driving force behind stock price movements. All these shocks can be expected to shift firms' marginal revenue productivity.

is likely to cause my estimates to be conservative. This follows from the fairly innocuous assumption that firms do not cut amenities (or workers don't reduce their common perception of firms' amenities) by too much on average when idiosyncratic productivity improves. This assumption implies my elasticity estimates are an upper bound on the true elasticity. My estimation strategy correctly identifies the supply elasticity when I further invoke the identifying assumption that only market-specific shocks move the level of the labor supply curve. This assumption has been imposed in essentially all papers that estimate passthrough parameters of firm-specific shocks to workers (see Lamadon et al. (2019), Kline et al. (2019), Garin and Silverio (2020), for example).[13] This is because unobservable firm-specific amenities shocks bias passthrough coefficients. This is an issue even when using quasi-experimental variation in firm revenue productivity, because one still cannot guarantee that amenities do not also move in response to the shock.

Even if there is some degree of firm-specific labor supply shocks, there are at least two reasons to prefer an instrumental variables estimate of supply elasticities over an ordinary least squares estimate. First, in internet appendix IA.1 I also examine the implied elasticity estimate from running an OLS regression. While it is easy to sign the bias for my IV estimates (which leans towards being conservative) for any volatility of firm-specific amenities/supply shocks, OLS elasticity estimates can be either upward or downward biased. Second, in practice there is likely at least some measurement error in firms' wages (for example, due to imperfections in the linkages between firm wage information in the LEHD with financial information in Compustat), which in practice would attenuate OLS estimates of elasticities towards zero and yield overly large wage markdowns.

In practice the economic magnitude of the bias in my estimates, if any, is likely to be small. In robustness checks detailed in section IA.3 of the internet appendix I examine elasticity estimates implied by several alternative instrument for shocks to revenue productivity; while each invokes slightly different identifying assumptions (which I cover in detail in the appendix), these alternative shocks yield quantitatively very similar elasticity estimates. I also identify sets of firms that were likely to have experienced observable firm-specific labor supply shocks, finding that my elasticity estimates are insensitive to the inclusion of these observable controls.

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

$$\log(Y_{j,t+1}) - \log(Y_{j,t}) = \alpha + \alpha_{I(j),t} + \beta \text{Stock Ret}_{j,t\to t+1} + \Gamma X_{j,t} + \epsilon_{j,t} \tag{5}$$

Here Y = either firm Compustat employment or firm average full-time full-year equivalent

[13] Papers in the asset pricing literature which incorporate labor market frictions (such as Kuehn et al. (2017), Belo et al. (2014), or 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.

adjusted wage.[14] One concern with Compustat employment is that it includes all employees, including those in non-domestic establishments; later I show that using employment from US administrative data in the Longitudinal Business Database yields quantitatively similar patterns in supply elasticity estimates, albeit somewhat smaller, meaning my choice to use Compustat employment yields conservative wage markdowns. The annual stock returns in (5) 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 $\widehat{\beta}^{Emp}/\widehat{\beta}^{Wage}$ gives the supply elasticity. The $\alpha_{I(j),t}$ denote 3-digit NAICS industry $\times$ year fixed effects, which control for common market shocks. The identifying assumption of the estimation strategy is that conditional on covariates, common productivity shocks (which in practice move the labor supply curve, as shown in internet appendix IA.1) operate at the market level, as defined by industry-year.

The controls $X_{j,t}$ 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. My ability to control for such composition changes derived from individual-level LEHD data is an improvement over prior firm-level analyses of labor demand responses to firm-specific shocks, such as Abowd and Lemieux (1993), or more recently Berger et al. (2021). 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. Because the regression requires growth rates in past and future wages and employment and lagged asset growth to all be observable, my analysis sample size is a bit smaller relative to my full matched sample at about 45,500 observations (rounded for Census disclosure purposes).

[14] The 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.

Before estimating elasticities via (5), I first look at local projections of wages and employment responses over different forward- and backward-looking horizons, which will reveal pre-trends in the relation ship 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:

$$\log\left(\frac{1}{|h|}\sum_{k=1}^{|h|} Y_{j,t+k\times\text{sign}(h)}\right) - \log(Y_{j,t}) = \alpha + \alpha_{I(j),t} + \beta_h \text{Stock Ret}_{j,t\to t+1} + \Gamma X_{j,t} + \epsilon_{j,t} \quad (6)$$

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 (6) for horizons $h = -5$ to $h = 5$ years. For $h = 1$ year the specification (6) is the same as (5). Note also that for $h = -1$ the coefficients are set to zero by the lagged growth controls, and that the $h = 0$ horizon is zero by definition. I plot the results in panel A 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 next subsection I show that stock returns also have a highly significant positive impact on contemporaneous growth rates in firm log value-added per worker, further supporting that stock returns immediately shift the firm's marginal product of labor and hence move labor demand.

2.3 Baseline Elasticity Estimates

I report estimates of equation (5) in panels A and B of Table 1. Focusing on the first column, I find an estimate of 0.12 for employment growth in panel A and 0.046 for wage growth in panel B. In panel C I take the ratio of the two estimates to arrive at my baseline pooled average supply elasticity estimate of about 2.5 (standard error of 0.3), found in the first column of Panel A in 1. This is well in line with the range of estimates found in previous research, although on the slightly smaller end.[15] This is likely because of my focus on publicly-traded firms, which are much larger than the typical firm in the US economy.

The next two columns of Table 1 explore estimates of (5), modified to focus on incumbent workers versus new recruits, where incumbent workers who have been at the firm since at least the previous year. I recalculate the worker skill composition controls and the growth

[15] Sokolova and Sorensen (2021) find the median supply elasticity estimate is around 1.7 in a meta-analysis. Bassier et al. (2020) point out that recent quasi-experimental evidence—such as Kroft et al. (2020), Cho (2018), and Dube, Manning, and Naidu (2018), and Caldwell and Oehlsen (2018)—has found supply elasticities between 2 and 5.

rates in the employment counts and wages separately for incumbents and recruits.[16] I find larger wage responses for incumbent workers and larger employment responses for recruits, which translates into a smaller supply elasticity for incumbents in panel C: 1.7 (standard error of 0.22) for incumbents versus 7.2 (standard error of 0.49) for recruits. As I show when I cover my model in section 3, incumbents also enjoy a wage premium over recruits of the same skill level (as defined by worker-level AKM intercept) of about 15%. These facts together are consistent with incumbent workers being expensive to replace (Jager and Heining, 2019; Kline et al., 2019), where the incumbent wage premium is a retention mechanism used to avoid the costly loss of existing employees. I show in my model in section 3 that this implies firms dial back their exercise of market power over these workers considerably relative to the static wage markdown formula coming from just the supply elasticity in equation (4). I quantify exactly how much this channel matters when I match the data moments from this section with my model.

In addition to my estimates by all workers and incumbents/recruits, I also 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 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 heterogeneity in 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 that this is not the case: despite strongly increasing valuations across labor productivity quartiles, there is no evidence that firms in

[16] Since I can't directly observe the share of Compustat employment by worker incumbent/recruit 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, $N^k_{j,t}$ of which are of type k, where k indexes incumbents or recruits, then I assume Compustat employment in group k is equal to $EMP^k_{j,t} = EMP_{j,t} \times N^k_{j,t}/N_{j,t}$. I continue to use Compustat employment because it is reported at year end, which allows more time for employment to respond to a stock return shock than LEHD employment. Compustat employment also reflects workers at all the firm's establishments instead of just those at the establishments that are covered by the LEHD in that year, which is helpful especially in the first part of my sample when the LEHD doesn't cover all states.

the top quartile of labor productivity have higher average investment rates from the firms in the bottom quartile. These stylized facts strongly point to empirical labor productivity dispersion reflecting dispersion in economic rents earned from market power in either product markets or labor markets.

With these characteristics of productive firms in mind, I estimate a modified version of (5) in the final four columns of Table 1 to test if differences in labor market power can speak to these differences 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.

$$\log(Y_{j,t+1}) - \log(Y_{j,t}) = \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} + \sum_{q=1}^{4} \mathbf{1}(q(i,t) = q) \times \beta_q \text{Stock Ret}_{j,t\to t+1} + \Gamma X_{j,t} + \epsilon_{j,t} \tag{7}$$

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 (5), 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.

The employment growth response estimates in the last four columns of panel A relative to the wage growth responses in panel B of Table 1 show that productive firms must increase wages by more to yield a given increase in employment, implying more inelastic labor supply, and hence larger wage markdowns from marginal product. In panel C I estimate a labor supply elasticity of 1.17 (standard error of 0.15) for firms in the top labor productivity quartile and 4.32 (standard error of 0.87) for firms in the bottom labor productivity quartile. The difference between these two elasticity estimates is highly significant, with a p-value of about 0.001. These estimates immediately imply a wide gap in rents earned from wage markdowns across the productivity distribution, which in turn can affect the differences in labor shares, valuations, and income to capital owners in the cross-section of firms sorted on labor productivity. In section 4 I combine my estimates from this section with my model calibration to quantify the implications of these estimates in terms of these average outcomes, finding a quantitatively substantial role for rents earned from labor market power in explaining the patterns documented in Table 2.

Besides labor supply elasticities, my approach allows me to estimate rent-sharing elasticities through a stock-return induced shock to value-added per worker. In panel D of Table 1 I estimate a modified version of (5), where the dependent variable is now the growth rate in a firm's value-added per worker, and I replace the one-year prior wage and employment growth controls with growth in value-added per worker. The first three columns of panel D covering the whole firm show the exact same regression, since it doesn't make sense to define rent-sharing elasticities with respect to value-added per incumbent or recruits. Panel D shows a highly significant response of value-added per worker growth to stock returns, with a coefficient of 0.18 (standard error of 0.017), confirming that stock returns are a powerful shifter for the firm's marginal product of labor. In panel E I compute the rent-sharing elasticity as the ratio of the wage estimates in panel B to the value-added per worker estimates in panel D. The pool average rent-sharing elasticity in the first column of panel E is 0.25; this masks heterogeneity by incumbent and recruit status in the next two columns, with incumbents having twice as large a rent-sharing elasticity (0.29 versus 0.15). This pattern is qualitatively consistent with Kline et al. (2019), though unlike them I do find a positive elasticity for recruits, which reflects some degree of imperfect competition for workers both within and outside of the firm.

In internet appendix IA.2 I show that both my modeling framework and a closely related multilateral bargaining model imply increasing rent-sharing elasticities as labor productivity increases. In the final four columns of panel D, I estimate a version of equation (7), where again I replace employment or wage growth with value-added per worker growth. Each labor productivity quantile exhibits a highly significant positive response of labor productivity growth to stock returns. I again take the ratio of the estimates in panel B to panel D, and report these in the last four columns of panel E. I find a monotonic increase in rent-sharing elasticities as labor productivity increases, with an estimate of 0.16 for firms in the bottom quartile of labor productivity and 0.35 for those in the top quartile. Despite a large literature estimating rent-sharing in various contexts (Card et al., 2013), I am the first to document such patterns of cross-sectional heterogeneity in the passthrough elasticities of value-added per worker to wages.[17] I include the 7 rent-sharing elasticities from panel E of Table 1 as untargeted moments for my model, and I find that the model delivers quantitatively accurate predictions for all these moments.

2.4 Time-Variation in Labor Supply Elasticities for Firms Sorted on Labor Productivity

To examine the role of my elasticity estimates in explaining time series changes in the labor share in my sample of Compustat firms, I re-estimate (7) 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 2. 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. This secular decrease in labor supply elasticities implies a role for changes in labor market power to explain time-series in the labor share, which I quantify explicitly in section 4.

[17] In prior work Chan et al. (2021) estimate heterogeneous TFP to wage passthrough elasticities, as opposed to rent-sharing elasticities, using Danish data.

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 the cross-sectional gap in implied markdowns widened between the first and second half of my sample: in the first half of the sample, productive firms' implied markdowns would have paid workers a share of about 26% less of their marginal product compared to unproductive firms in the 1991-2002 period (1.34/2.34 minus 4.86/5.86), whereas they would be expected to pay about 31% smaller share of marginal product in the 2003-2014 period (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 higer frequency. I estimate (7) 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

$$\text{Log Markdown Spread}_t = \log\left(\frac{\widehat{\epsilon_{4,t}}}{1+\widehat{\epsilon_{4,t}}}\right) - \log\left(\frac{\widehat{\epsilon_{1,t}}}{1+\widehat{\epsilon_{1,t}}}\right) \tag{8}$$

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:

$$\text{Log Lshare Spread}_t = (1/3)\sum_{\tau=t-1}^{t+1}\left(E[\log(lshare_{4,\tau})] - E[\log(lshare_{1,\tau})]\right) \tag{9}$$

Under a Cobb-Douglas production function the log labor share is linear in the log wage markdown, so if these estimates reflect genuine differences market power then they should move together over time.

Figure 3 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 3 echoes the changing relative valuations and productivity levels of high- and low-productivity firms in Figure IA.4, suggesting the phenomena may be jointly linked. In accordance with this, the model I introduce and calibrate in Section 3 forges a direct link between productivity advantages, labor market power, and firm labor shares.

My results in this section imply that labor market power is likely to be a quantitatively important determinant of firm rents and cashflows, especially for the most productive firms. While my estimates are informative about the magnitudes of labor market power, they may not be sufficient by themselves to quantify the cashflows 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 formalize this concept in section 3, where I propose a model to map my empirical reduced-form supply elasticity estimates into quantitative markdowns in the presence of labor market adjustment costs.

2.5 Robustness Checks

I run a host of robustness checks for my estimates of (5), which I cover in more detail in section IA.3 of the internet appendix. Since a major 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 IA.3.1. One may also worry about the quantitative impact of other sources of endogeneity. In internet appendix IA.3.2 I instead estimate supply elasticities with several alternative labor demand shifters—including customers' stock returns, patenting, and earnings announcement returns—finding very similar estimates. In IA.3.3 I also show that controlling for salient observable firm-specific labor supply shocks, such as unionization events or changes in non-compete enforceability, have essentially no effect on the average supply elasticity estimate. The common finding from these exercises is that my baseline average supply elasticity estimate is very close economically and statistically to the various alternatives, suggesting that 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.4. In appendix Table IA.9 I sort firms into labor productivity quartiles within 2-digit NAICS industry. In appendix Table IA.10 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 replacing 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 İmrohoroğlu 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.[18] 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 both this point).

Three tests confirm that differences in stock-based incentive pay are not driving my cross-sectional findings. In appendix Table IA.11 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. Finally, in appendix Table IA.12 I replace stock returns with the same alternative labor demand shocks as in appendix Table IA.8, except I now group firms into high- and low-productivity bins. Productive firms continue to face significantly lower elasticity estimates when using these alternative shocks to labor demand. For further details on these estimates and the rationale behind them, refer to internet appendix section IA.4.2.

3 Dynamic Wage Posting Model

In this section I introduce a dynamic wage posting model where a single firm hires from a fixed pool of potential employees. I intentionally make the model parsimonious, including just the key ingredients that directly speak to patterns in the data documented in Table 1. Since the model features a single firm, I use time-series productivity differences in the model to proxy for cross-sectional differences in productivity in the data.

The primary purpose of the model is to quantify how much adjustment costs affect markdowns from marginal product relative to the standard static elasticity-based markdown formula; in section 4 I combine the model calibration and my empirical estimates to quantify how much value firms derive from wage markdowns and implications for firm labor shares. The model is a dynamic extension of the static framework in Kline et al. (2019), and is also closely related to the setup of Card et al. (2018).

3.1 Setup

Consider the dynamic wage posting problem of a labor market with a single firm. There are a fixed total of L workers in the market, with a fraction λ_t currently employed at the firm at the start of period t. The firm can make different wage offers to outside hires ("O") and

[18] Note 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.

incumbents ("I"). The firm faces upward-sloping labor supply curves for these two classes of worker. For incumbents:

$$L_t^I(w_t^I) = \lambda_t L \left(\frac{w_t^I - b}{\overline{w} - b}\right)^\beta \tag{10}$$

For outside hires:

$$L_t^O(w_t^O) = (1 - \lambda_t) L \left(\frac{w_t^O - b}{\overline{w} - b}\right)^\beta \tag{11}$$

This labor supply functional form is borrowed from Card et al. (2018) and Kline et al. (2019). The term $\left(\frac{w_t^k - b}{\overline{w} - b}\right)^\beta$ represents the probability that a worker of type $k = I, O$ accepts the offer to work at the firm. In Kline et al. (2019), this derives from each worker receiving an unverifiable outside offer, with wage-equivalent distribution values drawn from a CDF $G(x) = \left(\frac{x - b}{\overline{w} - b}\right)^\beta$ with support in $(b, \overline{w}]$. Similarly, I restrict the support of wage offers to be between b and $\overline{w}$ to ensure that the probability a given worker chooses to be employed at the firm is bounded between 0 and 1. Card et al. (2018) derive the above labor supply function in a model of worker discrete choice with unobserved taste shocks between firms, where b is a lower bound wage offer at which no worker will choose to be employed. In Card et al. (2018) the parameter β is inversely related with the variance in workers' unobservable preferences among firms; higher values of β lead to more elastic labor supply. Like Kline et al. (2019) I take the functional form in equations (10) and (11) as given. Differently from them, I assume that firms choose the wage offer for both their incumbent workers and outside hires.

As documented in section IA.2 of the internet appendix, another nice feature of this functional form is that it leads to wage offers that are closely related to a bargaining model in the vein of Stole and Zwiebel (1996), where multiworker firms bargain with employees over the marginal surplus; $\beta/(1 + \beta)$ acts like the workers' bargaining weight and b functions like the workers' outside option. Unlike these bargaining frameworks, however, firms make wage offers to choose their optimal retention rate. This mechanism appears to be a prominent feature of the data and is directly implied by my wage posting framework, but would be hard to capture in the related bargaining alternatives where wages do not affect retention rates.

In order to focus on the firm's decision, I deliberately make the worker employment decision simple. An incumbent or potential recruit takes the wage offer at time t and decides whether or not to be employed at the firm for that period according to the choice probabilities implied in (10) and (11). Though β and b could in theory be worker-type specific, I restrict them to be the same for incumbent workers and outside hires in order to limit the number of free parameters when I calibrate the model.

The supply elasticity facing the firm for workers of type $k = I, O$ is

$$\frac{\beta w_t^k}{w_t^k - b} \tag{12}$$

This is decreasing in the wage and is increasing in β and b. In this setup, more productive firms will endogenously choose to inhabit a more inelastic portion of their labor supply curve.

Firms may pay outside hires differently because they are imperfect substitutes for incumbents. Specifically, hiring outside workers incurs quadratic adjustment costs:

$$c_t(L_t^O) = \exp(\sigma_c z_t)\bar{C}\lambda_t L \left(\frac{L_t^O}{\lambda_t L}\right)^2 \tag{13}$$

The adjustment costs may represent the cost of training workers in firm-specific production technologies or recruitment/hiring costs. The $\lambda_t L$ term gives the number of incumbent workers employed at the firm at the start of time t, so adjustment costs are quadratic in the ratio of outside hires to start of period incumbents (the "hiring rate"). I include the term $\exp(\sigma_c z_t)$ to allow adjustment costs for a given hiring rate to be increasing in firm productivity, since high labor productivity firms also tend to hire more skilled workers empirically (see appendix Table IA.2). Given that Belo, Li, Lin, and Zhao (2017) and Jager and Heining (2019) finding that adjustment costs are likely higher for skilled labor, this feature is a reduced form way of getting these costs to increase in firm "skill" in a setting where there is only one type of labor.

The firm produces output according to a Cobb-Douglas production function in the sum of inside and outside hires. Firm log productivity z_t follows an AR(1) process. There is no aggregate risk and firm cashflows are priced at the constant risk-free discount rate r. The firm solves

$$V_t = \max_{w_t^I, w_t^O} F_t(L_t^I, L_t^O) - \left(w_t^I L_t^I + w_t^O L_t^O + c_t(L_t^O)\right) + \exp(-r)E_t[V_{t+1}] \tag{14}$$

Subject to $L_t^O = (1-\lambda_t)L\lambda_t L\left(\frac{w_t^O - b}{\overline{w}-b}\right)^\beta$ and $L_t^I = \lambda_t L \lambda_t L\left(\frac{w_t^I - b}{\overline{w}-b}\right)^\beta$

$$\lambda_{t+1} = (L_t^O + L_t^I)/L \tag{15}$$

Where

$$F_t(L_t^I, L_t^O) = \exp(z_t)\left(\bar{z}(L_t^I + L_t^O)\right)^{(1-\alpha)} \tag{16}$$

$$z_{t+1} = \rho_z z_t + \sigma_z \epsilon_{z,t+1} \tag{17}$$

and $c_t(L_t^O)$ is the cost of hiring L_t^O workers. The law of motion for λ_t means that a time-t retained incumbent or outside hire becomes an incumbent at the start of $t+1$.

3.2 Wage Markdowns with Adjustment Costs

The presence of adjustment costs changes the markdown formula in the standard monopsony framework. To see this, we take the first-order conditions of (14) for worker types

$k = I, O$, resulting in a Q-theory like optimality condition for labor demand:

$$w_t^k = \frac{\varepsilon_k(w_t^k)}{\varepsilon_k(w_t^k)+1}\left(MRP_t^k - \frac{\partial c_t}{\partial L_t^k} + \exp(-r)E_t\left[\frac{\partial V_{t+1}}{\partial L_{t+1}^I}\right]\right) \tag{18}$$

Here MRP_t^k is $\frac{\partial F_t}{\partial L_t^k}$, which is the same for incumbents and recruits given the production function $F_t(L_t^I, L_t^O) = \exp(z_t)\left(\bar{z}(L_t^I + L_t^O)\right)^{(1-\alpha)}$. There are two additional terms in (18) relative to the static markdown formula. The first is $\frac{\partial c_t}{\partial L_t^k}$, the derivative of the adjustment cost function with respect to an additional unit of labor of type k. This term is zero for incumbents and positive for outside hires. The presence of adjustment costs leads the firm to prefer retaining incumbent workers over replacing them with outside hires, generating a wage premium for incumbent workers over recruits.

The second additional term relative to a standard static wage markdown in the dynamic wage offer equation is $\exp(-r)E_t\left[\frac{\partial V_{t+1}}{\partial L_{t+1}^I}\right]$, which is the time-$t$ expected marginal benefit of having an additional incumbent worker at time $t+1$. Because both a retained incumbent or recruit hired today becomes an additional worker in the available pool of incumbents tomorrow, this term is the same for both types of workers. This term is also always positive because entering the next period with more incumbents is strictly better than having fewer incumbents. This implies that wage markdowns from marginal product will be strictly smaller for incumbent workers than would be inferred by supply elasticities alone, and may be larger or smaller for recruits depending on the relative size of $\frac{\partial c_t}{\partial L_t^O}$ and $\exp(-r)E_t\left[\frac{\partial V_{t+1}}{\partial L_{t+1}^I}\right]$. Note also that since incumbents and outside hires face different wage offers, the supply elasticities themselves also differ because the elasticity (12) is a decreasing function of the wage offer for each worker type.

Define for worker types $k = I, O$

$$\textit{dynamic markdown}_t^k \equiv \frac{w_t^k}{MRP_t^k - \frac{\partial c_t}{\partial L_t^k} + \exp(-r)E_t\left[\frac{\partial V_{t+1}}{\partial L_{t+1}^I}\right]} = \frac{\varepsilon_k(w_t^k)}{\varepsilon_k(w_t^k)+1} \tag{19}$$

$$\textit{MRP markdown}_t^k \equiv \frac{w_t^k}{MRP_t^k} \tag{20}$$

$$\textit{markdown wedge}_t^k \equiv \frac{\textit{MRP markdown}_t^k}{\textit{dynamic markdown}_t^k} \tag{21}$$

The markdown expressions (19), (20), and (21), are worker type specific, so I introduce their firm-level aggregates:

$$\textit{dynamic markdown}_t \equiv \frac{\sum_k w_t^k L_t^k}{\sum_k \left(MRP_t^k - \frac{\partial c_t}{\partial L_t^k} + \exp(-r)E_t\left[\frac{\partial V_{t+1}}{\partial L_{t+1}^I}\right]\right) L_t^k} \tag{22}$$

$$MRP\ markdown_t \equiv \frac{\sum_k w_t^k L_t^k}{\sum_k MRP_t^k L_t^k} \tag{23}$$

$$markdown\ wedge_t \equiv \nu_t \equiv \frac{MRP\ markdown_t}{dynamic\ markdown_t} \tag{24}$$

I examine the properties of these objects in the calibration. The firm-level markdown wedge (24) is a particularly important quantity, which maps reduced-form empirical labor supply elasticity estimates into model-implied markdowns from marginal product.

3.3 Calibration and Model Fit

Closed form solutions for the firm's objective function (14) are not readily available, so I solve the model for a given calibration through policy function iteration. I calibrate the model to match the average separations rate; the incumbent wage premium; overall supply elasticity, incumbent/recruit specific supply elasticities; the relative stock-return wage passthroughs of incumbents and recruits; and the average log labor share.[19] I estimate model-implied supply elasticities by running a regression of model-generated employment and wage growth on stock returns (growth in the value function). I use the Rouwenhourst (1995) method to create a discretized grid of 9 points to approximate the AR(1) productivity process. For pre-calibrated parameters, I normalize the labor force size L to 1; set the annual discount rate to .02; and set $\overline{w}$ to 1; I take $\rho_z = 0.9$ to roughly match the persistence of annual log value-added per worker in the data; and finally, I set both the baseline labor productivity $\bar{z}$ and convex adjustment cost volatility σ_c equal to 1.

Panel A of Table 3 shows the model calibration of the externally and internally calibrated parameters, and Panel B of Table 3 displays the model fit to targeted moments. There are 5 internally calibrated parameters for 7 targeted moments. The quantitative fit is very good, with almost exact matches on average incumbent wage premium, separations rates, relative wage passthroughs, and average log labor shares; the model calibration does slightly underestimate average labor supply elasticities, but is still pretty close quantitatively. I calibrate the parameters $\beta = 0.38$, $b = 0.545$, and $\sigma_z = 0.145$ to help match the average supply elasticities and the incumbent/recruit stock return wage passthrough ratio; the parameter $\bar{C} = 0.55$ is identified primarily by the incumbent-recruit wage premium. The parameter $\alpha = 0.22$, which is one minus labor returns to scale, and productivity volatility $\sigma_z =$ jointly help me get close to the average empirical supply elasticity and log labor share. Finally, the average wage levels implied by the combination of parameters help to jointly match the empirical average firm-level separations rate.

While my model parameters don't always have the exact same structural interpretation, I next compare them to some reasonably close analogues from recent prior literature to show

[19] The incumbent wage premium is skill-adjusted and represents the intercept from a regression of the wage premium on the incumbent minus recruit difference in average AKM worker effects.

that my calibration seems to be in a reasonable range relative to related parameters in prior research. I look at the calibrations of the dynamic models in two other papers, Kuehn et al. (2017) and Belo et al. (2017), which both have model features that are related to mine. Kuehn et al. (2017) calibrate the benefit of unemployment to be 0.71; my analogous parameter is the outside option of b, calibrated at 0.545. Kuehn et al. (2017) also have a worker bargain weight of 0.11. The most comparable parameter in my calibration is $\beta/(1+\beta) = 0.27$, which is a little higher. However, the parameter β also affects the supply elasticity, and so $\beta/(1+\beta)$ doesn't have the exact same structural interpretation as in a pure bargaining model, since it it also directly affects the endogenous labor supply response to wage changes. Belo et al. (2014) calibrate the high-skill convex adjustment cost parameter to 1.8 and the low-skill convex adjustment cost parameter to 0.17; my calibration of $\bar{C} = 0.55$ lies in between the two. Finally, my calibration of returns to scale of $1 - \alpha = 0.78$ lies in between those of Kuehn et al. (2017) and Belo et al. (2017) (0.75 and 0.85, respectively).

A supply elasticity that decreases in the wage offer is a key feature of the model that helps match the data moments. Consistent with the data, recruits face higher average supply elasticities. In the model this is because their wage offers are closer to the reservation wage level b, and elasticities can be arbitrarily high as wage offers approach b. This feature of the labor supply function also implies that firms face lower elasticities in more productive states. In Panel C of Table 3 I show that the model is matches the monotonic pattern in supply elasticities and log labor shares when sorting by labor productivity without explicitly targeting these moments. Interestingly, although the model matches the monotonic pattern in firm labor shares it misses on the magnitude of the spread. The log labor share spread is 0.37 in the model but 0.72 in the data. In the next section I show that this is also the case for my empirical supply elasticity estimates–although they imply a spread in labor shares, the spread is wider in the data than implied by my empirical estimates. This motivates a decomposition of the labor share spread in Section 4, where I examine how much of a role there is left over for differences in markdowns or labor returns to scale after accounting for the implied spread in markdowns. Including additional model ingredients, such as product market power that also increases with firm productivity, could likely explain the remainder of this empirical spread.

In panel D of Table 3 I show the model fit to the heterogeneous rent-sharing elasticities that I estimated in empirically in the last panel of Table 1, representing 7 additional untargeted moments for the model. The model delivers an overall rent-sharing elasticity of 0.30, which is close to the 0.25 estimate in the data; the model also captures the difference in incumbents versus recruits, with an elasticity of 0.36 for incumbents in 0.19 for recruits (compared to 0.29 and 0.15 for their data counterparts). Finally, the model does quite well at matching the slope of rent-sharing elasticities across the productivity distribution, exactly matching the rent-sharing elasticity of 0.15 in the bottom quartile, and delivering a rent-sharing elaticity of 0.38 for the top quartile compared to 0.35 in the data. In total, panels C and D represent 15 untargeted moments in the data on top of 7 targeted moments to which the model fit is impressively accurate despite only 5 free parameters. Thus the model is remarkably successful in simultaneously capturing these rich patterns of heterogeneity in the data.

I now analyze the behavior of markdowns in the model. In the left panel of Figure 4 I compare average dynamic markdowns implied by estimating supply elasticities in a regression of employment and wage growth on stock returns versus the actual average dynamic markdown from the true supply elasticities determined by (12). Because of adjustment costs, the regression-implied elasticities slightly underestimate the wages paid to incumbents and overestimate the wages paid to recruits. However, when aggregating to the firm level, the linear regression elasticity-implied markdown closely reflects the average dynamic markdowns for the whole firm. In the second panel of Figure 4 I compare the model dynamic markdowns to the marginal revenue product markdowns. This comparison shows the difference between the markdowns that would be implied by using the static expression for markdowns based on supply elasticities alone versus the actual model markdown from marginal product. At the firm level, the average dynamic markdown is 0.68, which implies that firm pays about 68% of the denominator in (22) to its workers in wages on average. However, this again masks heterogeneity by worker type. Because incumbents are costly to replace, their wages are only marked down to 83% from marginal product; without accounting for dynamics due to adjustment costs, supply elasticities alone would imply a wage of 64% of marginal product. The reverse is true for recruits: supply elasticities would imply a wage that is about 80% of marginal product, but recruits' wages are actually about 71% of marginal product.

A key object for empirical quantification is the average firm-level markdown wedge implied by the model. This is about $0.79/0.68 \approx 1.16$, so wage markdowns at the firm level are about 16% smaller in magnitude after accounting for adjustment costs than would be inferred by the simple static markdown formula $\epsilon/(1+\epsilon)$. I label the markdown wedge ν. In Table 4 I show the markdowns from marginal revenue product that result from my supply elasticity estimates in section 2 combined with the markdown wedge ν.[20] Even after adjusting for dynamics due to adjustment costs, the markdowns still imply a considerable amount of monopsony power. However, adjustment costs do negate labor market power quite a bit. The baseline empirical elasticity estimate of 2.52 yields a markdown that is 72% of marginal product; after adjusting for dynamics this becomes 83% of marginal product.[21] A supply elasticity of about 4.9 would give this markdown in a static setting. Table 4 combines the model markdown wedge with my empirical elasticity estimates. The table also shows that while low productivity firms do

[20]For simplicity I report results assuming the markdown wedge is constant across productivity levels. Results are very similar quantitatively if I allow the markdown wedge to vary by productivity.

[21]Interestingly, despite being derived from very different methods, both of these numbers are very similar to the estimates that Chan, Mattana, Salgado, and Xu (2023) derive using production function estimation techniques to back out the impact of wage markdowns and adjustment costs in Denmark.

mark wages down, they still pay 94% of total marginal product. This is reasonably close to a competitive benchmark. On the other hand, the most productive firms pay 62% of marginal product even after adjusting using the markdown wedge ν.

3.3.1 Calibration for 1991-2002 and 2003-2014 Subperiods

Given the downward trends in supply elasticities and estimated markdown spreads that I find in Figures 2 and 3, I introduce calibrations to match moments for the first and second halves of the sample (1991-2002 and 2003-2014, respectively). The calibrated parameters by subperiod are in appendix Table IA.5. The average cross-sectional standard deviation of log value added per worker increases by about 5% over the two subperiods, so I decrease σ, the productivity dispersion parameter, by about 2.5% relative to the main calibration for the 1991-2002 period and increase it by 2.5% for the 2003-2014 calibration. I calibrate the remaining parameters to match the same moments as in my main calibration.

In Table IA.6 I give the corresponding model and data moments for the subperiods. Supply elasticities and separations rates decline in the data across the two periods; I decrease the parameter β, which governs the average supply elasticity, and b, which governs the worker outside options to match these moments. My calibration of β decreases from 0.5 to 0.34 from the first half to the second half of the sample, and b decreases from 0.549 to 0.515. This agrees in principle with Stansbury and Summers (2020), who suggest that worker power in the workplace has declined over recent decades.

The incumbent log wage premium increased by about 30% (from 0.13 to 0.17) across the two periods. This suggests increased adjustment costs, and I increase $\bar{C}$, the convex adjustment cost parameter, from 0.38 to 0.745 to match these moments. The resulting model-implied markdown wedge ν increases from 1.11 to 1.24 between the first and second halves of the sample period, mitigating some, but not all, of the impact of decreasing supply elasticities. I use these subperiod-specific elasticity estimates and markdown wedges when I quantify the model in the next section.

4 Quantifying Empirical Estimates

In this section I use the model-implied markdown adjustment parameter ν in conjunction with my empirical elasticity estimates from Section 2 to quantify the value that firms derive from their wage markdowns. I also use the model and estimated elasticities to back out the share of variation in labor shares across time and between firms of different productivity types that can be ascribed to differences in markdowns; finally, I briefly examine the extent to which wage markdowns can explain variation in average valuation ratios when sorting on labor productivity.

4.1 Value of Cashflows Derived from Labor Market Power

With the markdown wedges ν from the model and the empirical elasticity estimates for firms sorted on labor productivity, I now examine the counterfactual cashflows firms would bring in if they were not able to mark down wages from marginal product, but still held their production decisions constant. This is different from the competitive counterfactual, because in the competitive equilibrium both quantities and prices of labor would adjust. Instead, this counterfactual transfers the capital income attributable to wage markdowns from the actual owners of the firm's capital to the firm's workers after production has occurred. I use two measures of cashflows. My primary measure is based off the firm's operating income. I use the Compustat variable OIBDP adjusted for changes inventories, which was also used as the basis for my computation of firm value-added in (A.1). Hartman-Glaser et al. (2019) use OIBDP as a proxy for capital income available to the owners of debt and equity issued by the firm. The ratio of the value of wage markdowns to operating income can then be thought of as representing the share of capital income coming from wage markdowns. Going forward I use operating income and capital income interchangeably.

Denote $\widehat{\epsilon_{q,t}}$ the elasticity estimate for productivity quartile q at time t, where I use the elasticity estimates for either the 1991-2002 or 2003-2014 subperiods depending on which period year t falls in. Let $q(j,t)$ give the productivity quartile of firm j at time t. Finally, let ν_t denote the markdown wedge estimated over the subperiod where year t falls. Counterfactual labor expenses without markdowns are

$$\widetilde{LABEX}_{j,t} = \nu_t^{-1} \frac{\widehat{\epsilon_{q(j,t),\, t}} + 1}{\widehat{\epsilon_{q(j,t),\, t}}} LABEX_{j,t} \tag{25}$$

And the resulting markdown share of operating income:

$$\text{Markdown Share}_{j,t}^{OI} = \frac{\widetilde{LABEX}_{j,t} - LABEX_{j,t}}{OI_{j,t}} \tag{26}$$

I then compute the average markdown share of operating income:

$$\text{Average Markdown Share of OI} = \frac{1}{N} \sum_t \sum_j \text{Markdown Share}_{j,t}^{OI} \tag{27}$$

Here N denotes the total number of firm-year observations in the sample. Some firms report negative operating income, and occasional extreme values of Markdown Share$_{j,t}^{OI}$ skew the firm-level averages, so in (27) I focus on the subset of firms for which $OI_{j,t} > \widetilde{LABEX}_{j,t} - LABEX_{j,t}$. Alternatively, I report the median fraction of operating income derived from wage markdowns among a all firms.

Finally, I look at the aggregate share of operating income that comes from wage markdowns by summing true and counterfactual income across the whole population of firms:

$$\text{Aggregate Markdown Share of OI} = \frac{1}{T}\sum_t \left[\frac{\sum_j \widetilde{LABEX}_{j,t} - LABEX_{j,t}}{\sum_j OI_{j,t}}\right] \tag{28}$$

with T the number of years in the sample. I compute (27) and (28) for all firms and separately for firms in the top and bottom quartiles of labor productivity.

I report these averages in Table 5 overall and for firms in the top and bottom quartiles of labor productivity. In Panel A I focus on results for the full sample period. The first two rows of Panel A show the firm-level mean and median shares of operating income derived from paying wages that are different from marginal product. The average firm in the sample earns about 34% of its operating income by paying wages lower than marginal product; this figure is 30% for the median firm. There is large heterogeneity between the most- and least productive firms, with those in the bottom quartile of productivity earning only 17% of their income from wage markdowns, while firms in the top quartile earn about 43%. Since the operating income represents earnings available to the firm's capital owners, these estimates suggest that about a third of the capital income generated by the typical firm comes from wage markdowns.

In the third row of Panel A in Table 5 I compute the aggregate operating income share of wage markdowns, as in equation (28). Because markdowns are concentrated amongst larger firms, the aggregate average is higher than the firm-level average, at about 40%. Hence two-fifths of the capital income generated by publicly-traded firms were attributable to wage markdowns over the 1991-2014 period. Assuming these cashflows would be discounted at the same rate as capital income overall, this has the interpretation that the dollar value of wage markdowns at the monopsony equilibrium is worth roughly 40% of the total enterprise value of publicly-traded firms. This different than saying the aggregate value of publicly-traded firms would be 40% lower without wage markdowns, because the exercise here holds equilibrium labor demand constant. The decline in total enterprise value in the counterfactual competitive equilibrium would be substantially smaller than 40%. I discuss this point further later on in this section.

Panel B of Table 5 breaks down these figures by the 1991-2002 and 2003-2014 subperiods. At the firm level, the average markdown share of operating income increases slightly from 0.34 to 0.35 between the 1991-2002 and 2003-2014. The median increases slightly more, from 0.29 to 0.32; finally, the aggregate operating income share also increases from 0.39 to 0.41. These slight overall increases mask heterogeneity by productivity type: the mean, median, and aggregate average share of operating income generated from wage markdowns actually dropped for the least productive firms, while all three rose for the most productive firms. The largest increase was for the median high productivity firm, which saw their wage markdowns as a share of operating income increase from 0.45 to 0.52 between the two periods. This divergence parallels the increase in productivity dispersion between top and bottom firms over the sample period, as shown for example in Figure 3.

4.1.1 Interpreting the Magnitude of cashflows Generated from Wage Markdowns

Given the size of wage markdowns as a fraction of operating income in Table 5, it is helpful to provide some context for interpretation and to compare the magnitude against other results.

First of all, note that these 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 40% lower if we imposed perfect competition. To illustrate, in section IA.6 of the internet appendix I make a back of the envelope calculation in a simple static representative firm setting where wage markdowns are calibrated to be exactly 40% of operating income, and the production function is calibrated to roughly match the average aggregate labor share. In this setup, firm value declines by 16% in the counterfactual competitive equilibrium relative to the monopsony equilibrium.

How does the magnitude of my estimates compare with other estimates of market power? In the next subsection I introduce two measures of aggregate labor shares; depending on the measure the mean aggregate labor share is about 49-55% of value-added, which implies that wage markdowns average about 18-20% of total output. In comparison, 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, so my estimates are about half as large as this figure. The De Loecker et al. (2020) estimates are also closely related to the ratio of sales to cost of goods sold, and they assume there is no monopsony power. Imperfect competition in labor markets can also affect the ratio of sales to costs of goods sold, so their estimates likely reflect both labor and product market power. Yeh et al. (2022) jointly estimate wage markdowns and price markups for manufacturing firms using a production function approach that is related to the De Loecker et al. (2020) method. They find that the average firm pays about 65% of marginal product. Since they estimate that markdowns increase in firm size, the value-weighted gap would be even wider. In comparison, I find that total aggregate wages are about 70% of aggregate marginal product, which is smaller but of a comparable magnitude.

4.2 Decomposition of Labor Share Differences in the Cross-Section and Time Series

Next, I look at what fraction of the gap in labor shares between high- and low-productivity firms can be explained by differences in wage markdowns; I also examine how changes in the aggregate labor share in the time series were impacted by changing markdowns. In this section I use two measures of labor shares. My baseline, defined previously in (A.4), comes from imputed labor expenses $LABEX_{j,t}$ derived from LEHD earnings and Compustat employment data. I find that $LABEX_{j,t}$ is on average a little lower than Compustat reported staff and labor expense (Compustat variable XLR), but the variable XLR is only available for a very small fraction of firms. To get around this I create an imputed version of XLR, which uses the fact that $LABEX_{j,t}$ and $XLR_{j,t}$ are very highly correlated when $XLR_{j,t}$ is observable. Details on the imputation are in appendix section A.3. I label this imputed version $\widehat{XLR}_{j,t}$, and create my alternative labor share measure by simply replacing $LABEX_{j,t}$ with $\widehat{XLR}_{j,t}$ in equations (A.1) and (A.4). I call this measure LShare $(\widehat{XLR})_t$. The logs of the two firm-level labor share measures share a high correlation of 0.91, though the time series behavior of the aggregate labor share from summing the two versions of labor expenses is a bit different; I elaborate on this point when I do the decomposition of aggregate labor share changes in the time series.

Table 3 shows that the model generates a monotonically decreasing pattern in labor shares as productivity improves, as is also found in the data. Still, the labor share spread in the model is not as wide as the empirical spread. This is also true when using the empirically estimated markdowns in Table 4, and leaves room for other factors, such as differences in product market power or production technologies, to explain the remainder of the labor share spread. For example, assume a Cobb-Douglas production function in labor, $F(L) = AL^{1-\alpha}$, but also suppose the firm chooses output according to inverse demand $P(Q) = P_0 Q^{-1/\gamma}$, where γ is the price elasticity of demand. Also assume that because of labor market power wages are a markdown μ from the marginal revenue product of labor. Then the revenue function is $R(L) = P_0 A^{1-1/\gamma} L^{(1-1/\gamma)(1-\alpha)}$, where $1 - 1/\gamma$ gives the inverse price markup. The log labor share satisfies

$$\begin{aligned}\log\left(\text{Labor Share}\right) &= \log\left(\frac{\mu R'(L)L}{R(L)}\right) = \log\left(\mu\right) + \log\left(1-1/\gamma\right) + \log\left(1-\alpha\right) \\ &= \log(\text{wage markdown}) - \log(\text{price markup}) + \log(\text{labor returns to scale})\end{aligned} \tag{29}$$

The difference in average log labor shares can then be ascribed to differences in average log wage markdowns, price markups, and returns to scale. What percentage of the average labor share differences between high- and low productivity firms can be attributed to differential wage markdowns? In Panel A of Table 6 I answer this question using the wage markdown estimates from Table 3 and my two measures of firm labor shares. I show estimated log markdowns and average log labor shares. The difference in empirical log markdowns from Table 4 is $\log(0.62) - \log(0.94) \approx -.41$, which I report in the bottom row of the third column of Panel A. The differences in the average log labor shares between high and low productivity firms if -0.72 for my baseline and -0.65 for the $\widehat{XLR}$ based measure; markdown differences can account for roughly three-fifths of the difference in average log labor shares (57% for my baseline labor share measure and 63.5% for the alternate measure).

My findings still leave an important quantitative role for product market power to play in fully explaining the cross-sectional gap in labor shares. If labor returns to scale are approximately the same between the two groups of firms then the remaining two-fifths could be explained by differences in price markups. The increasing product market power among dominant firms found by Autor et al. (2020) and De Loecker et al. (2020) and the widening markdowns that I find are likely to be jointly important phenomena quantitatively. Also note that my model implicitly forges a link between product and labor market power, because increases in product market power increase labor productivity, which in turn increases labor market power, both in my model and in the data.

I next compute the aggregate labor share using either $\widehat{XLR}$ or $LABEX$. Using $LABEX$, the aggregate labor share is given by

$$\text{Agg LShare } (LABEX)_t = \frac{\sum_j LABEX_{j,t}}{\sum_j VA_{j,t}} \tag{30}$$

and Agg LShare $(\widehat{XLR})_t$ is defined analogously. Here $VA_{j,t}$ is the value-added for firm j at time t, computed by adding either $LABEX_{j,t}$ or $\widehat{XLR}_{j,t}$ to operating income adjusted for changes in inventory as in (A.1). Although the firm-level variation in log labor shares for the two measures is very highly correlated, appendix Figure IA.2 shows that the downward trend in labor shares is more pronounced for Agg LShare $(\widehat{XLR})$. The average aggregate labor share drops from 0.579 to 0.524, when comparing the 1991-2002 and 2003-2014 subperiods; meanwhile Agg LShare $(LABEX)$ shows a more modest decline of 0.506 to 0.485.

Since the trend in Agg LShare $(\widehat{XLR})$ is much closer to the Compustat labor share trends found by Hartman-Glaser et al. (2019), in the next exercise I focus primarily on changes in Agg LShare$(\widehat{XLR})$, but show results for both versions. I compute the counterfactual labor share if markdowns remained at their 1991-2002 levels in the years 2003-2014. In particular, let $\eta_q^{1991-2002}$ denote the estimated wage markdown for productivity quartile q for the 1991-2002 period, and $\eta_q^{2003-2014}$ the estimated wage markdown for 2003-2014. The counterfactual labor share is given by

$$\widetilde{\text{Agg LShare}}\ (\widehat{XLR}) = \frac{\sum_j \widehat{XLR}_{j,t} \times \eta_{q(j,t)}^{1991-2002} \times \left(\eta_{q(j,t)}^{2003-2014}\right)^{-1}}{\sum_j VA_{j,t}} \tag{31}$$

It's important to be clear about what this counterfactual represents and what it does not. Similar to my prior exercise in section 4.1, this counterfactual asks the question: "Holding the distribution of workers across firms constant, how would the observed aggregate labor share change in the 2003-2014 period if enough final output was reallocated to workers to keep the gap between wages and marginal products constant at 1991-2002 levels?"

I compute Agg LShare $(LABEX)_t$ analogously. I then compute the counterfactual log change in the average labor share across the two periods. Panel B of Table 6 shows the results. In the first two columns I show the actual average labor shares for the two measures, and in the third column (labeled $2003-2014$) I give the counterfactual aggregate labor share. The first row shows that if markdowns had held constant at their 1991-2002 levels in the latter period, the average labor share Agg LShare $(\widehat{XLR})_t$ would have only declined from 0.579 to 0.554 instead of the observed level of 0.524. The actual log change, in the column Δ log(Lshare), was -0.10, whereas the counterfactual log change (Δ log(Lshare)) would have been -.044. This implies that about 56% percent of the decline ($\frac{0.10-0.044}{0.10} \approx 0.558$) can be explained by the change in wage markdowns. In the second row of panel B I compute the same counterfactual for Agg LShare $(LABEX)_t$. Because the labor share decline from this measure is less stark to begin with, the counterfactual actual predicts an *increase* in the labor share of about 1.8% instead of the observed 4.2% decrease.

In either case, I estimate that the change in wage markdowns can explain the bulk of the decline in average labor shares over the two periods. The results using my preferred measure Agg LShare $(\widehat{XLR})_t$ suggest that there is still some room for other proposed factors, such as increased price markups or labor-substituting technological change, to explain the remaining 44%.

4.3 Explaining Differences in Average Valuation Ratios by Labor Productivity

In Table 2 I showed that productive firms have high valuation ratios but don't have high investment rates, which could potentially be explained by differences in economic rents from labor market power pushing up valuations without reflecting any investment incentives. I now quantitatively examine to what extent wage markdowns can explain these differences in average valuation ratios. I again use total Tobin's Q from Peters and Taylor (2017), and I utilize a decomposition of valuation ratios from Crouzet and Eberly (2021) to predict the component that comes from cashflows derived from economic rents from labor market power. I then compute the counterfactual average valuations for firms of different productivity levels if these rents are taken away; see section A.4 of the internet appendix for implementation details of this empirical exercise.

Table 7 shows that the cashflows coming from wage markdowns can explain about two-thirds of the gap in average firm valuations between high- and low-productivity firms, and about 85% of the variance in average valuation ratios across all four productivity quartiles. This channel in turn explains why productive firms' investment rates are comparatively modest relative to their high average valuations.

5 Discussion

In this section I discuss potential causes of the rise in labor market power, economic and policy implications of my findings, and how my results can be distinguished with alternative explanations.

5.1 What Has Caused the Rise In Labor Market Power?

Two key features characterize the rise in labor market power. Rising productivity dispersion—illustrated in Figure IA.4 and found in numerous other papers[22]—has led highly productive firms to mark wages down by more relative to unproductive firms over time. Meanwhile, Figure 2 shows a generalized decrease in the supply elasticities of firms of all productivity levels between the 1991-2002 and 2003-2014 periods.

Declining elasticities suggest a secular increase in labor market frictions, so that workers have found it more difficult to substitute away from particular employers. Consistent with this, Molloy, Smith, Trezzi, and Wozniak (2016) find that the US labor market has become less dynamic over the past several decades, with less frequent worker transitions between firms. Similarly, in appendix Table IA.6 I show that the average firm-level separations rate has declined. My model calibration for the 1991-2002 and 2003-2014 periods suggests reduced outside options (model parameter b) and bargaining power (parameter β).

However, Berger et al. (2021), Rinz (2018) and Lipsius (2018) all find that local labor market concentration has been decreasing, which initially seems at odds with increasing aggregate labor market power. At the same time, many firms set relatively uniform wages nationally across their establishments because they view themselves as competing in a national labor market (Hazell, Patterson, Sarsons, and Taska, 2022). The importance of more national labor markets for skilled workers could also help explain in part why direct estimates of monopsony power have increased, even as local labor market concentration has decreased. For example, although skilled workers have lower rates of separation, Malamud and Wozniak (2012) and Amior (2020) document that conditional on moving, more educated workers make job moves that are far more geographically distant on average, and Diamond (2016) shows that educated workers tend to live in larger metropolitan areas which naturally have less local concentration. Also supporting a human capital specificity mechanism, Nimczik (2020) finds that more skilled workers tend to be employed in labor markets that are geographically dispersed but more concentrated in a set of particular industries.

[22]See Decker, Haltiwanger, Jarmin, and Miranda (2020), Autor et al. (2020), Hartman-Glaser et al. (2019), De Loecker et al. (2020), Kehrig and Vincent (2021), and Gouin-Bonenfant (2020) for example.

I briefly explore the quantitative role for workers of different skill levels in driving the pattersn I find by examining heterogeneity in elasticity estimates by worker skill. I group individual workers into low-, middle-, and high-skill workers—defined as in being in the bottom two quintiles, third and fourth quintile, or top quintile in the cross-sectional distribution of worker-level intercepts in the AKM wage decomposition—and I re-estimate supply elasticities via equation (5) by worker skill for overlapping moving 3 year windows to see which workers drive the declining supply elasticity trend. For comparison, I also do so for all workers. Appendix Figure IA.3 shows the results. While the supply elasticities for all skill levels decline over time, the most skilled workers have both the lowest elasticities overall, and are the most biggest drivers quantitatively for all workers. Thus there appears to be a drop in the supply elasticity across the skill distribution, the skilled workers appear to play an especially important quantitative role.

In Table IA.3 I re-estimate (7) for workers of the three skill groups to look at the spread in supply elasticities in the cross-section by labor productivity rank. In the top panel I show the estimates for all workers for comparison, and in the next three panels I show the breakdown by low-, middle, and high-skill workers, respectively. While labor supply elasticities decline in the cross-section as labor productivity increases for all skill groups, again I find that the skilled workers are again quantitatively important for driving the overall cross-sectional patterns. Additionally, the highly productive firms, which I find have the most labor market power also tend to hire more skilled workers on average (see appendix Table IA.2). This suggests that increased importance of skill specificity may be playing a role for skilled workers.

What could explain a rising importance of skill specificity? Skill-biased technological change has induced an increasing share of individuals to delay entry into the workforce in favor of obtaining higher education, potentially resulting in more specialized human capital in the process. Technological forces which have made large firms more dominant in an increasingly intangible-driven economy (see De Ridder, 2019, for example) may have also induced more skilled workers to be more intensively employed at dominant firms. While skilled workers may benefit from the high wages such employers can offer, these forces can reduce cross-firm labor substitutability and contribute to productive firms' widening labor market advantages. Song, Price, Guvenen, Bloom, and von Wachter (2018) find that the 1990-2015 period has seen increased sorting of productive workers to high-wage firms; appendix Figure IA.5 relatedly shows that the spread in average worker skill between productive and unproductive firms widened in the latter half of my sample period. Appendix Table IA.6 shows that the wage premium for incumbent workers has increased, suggesting that firms have found it more difficult to substitute workers from those outside the firm. All these patterns are consistent with specific human capital rising in importance.

Increased human capital specificity potentially has two opposing effects on the extent of labor market power: 1) employers find it more costly to replace their workers, which tends to raise the wages of incumbent workers and reduce monopsony power; and 2) workers are less able (or less willing) to substitute their labor across particular employers, which tends to increase monopsony power. My model calibration and empirical estimates suggest that the second effect has clearly dominated, especially at the top of the productivity distribution.

The overall rise of non-compete agreements, which now affect a large portion of the labor force and are focused on skilled workers (Starr, Prescott, and Bishara, 2021), could also have played a major part in reducing supply elasticities. These appear to be successful in reducing worker mobility and depressing wages, especially among skilled workers.[23] The time series change in labor market frictions also has interesting implications for the rise in wage inequality and the polarization of the US labor market (Acemoglu and Autor, 2011; Autor and Dorn, 2013). Since I find the lowest and most decreasing supply elasticities for skilled workers, this implies an even larger role for technological change in increasing wage inequality over the same time period.[24]

In all of the above I caution that more follow on work needs to be done to specifically delineate which labor market frictions are binding for which types of workers. My model, while providing an excellent fit to the data, is insufficient to specifically highlight which frictions workers of different skill levels face in labor markets in practice. There is, however, some initial supporting evidence for labor market power being driven by skill specificity: Goolsbee and Syverson (2023) find universities have substantial monopsony power over the tenure-track faculty; Prager and Schmitt (2021) find that hospital mergers affect wages more negatively when workers have specific skills; and, Lehr (2023) finds that the most successful innovative firms exert more market power over their inventors.

5.2 Economic and Policy Implications of Labor Market Power

Despite having low labor shares, productive firms pay far higher wages and have low worker separations rates (see appendix Table IA.4). This implies an important subtlety to consider for policy interventions intended to curtail labor market power: productive firms are likely to be coveted employers, which may explain much of their labor market competitive advantage. It has been well-understood that there are dynamic welfare tradeoffs in product markets between incentivizing ex-ante investment by allowing ex-post rents which may be Pareto inefficient in a static sense (Aghion, Bloom, Blundell, Griffith, and Howitt, 2005; Aghion, Bergeaud, Boppart, Klenow, and Li, 2022). The fact that successful firms exercise

[23]See Garmaise (2011), Balasubramanian, Chang, Sakakibara, Sivadasan, and Starr (2020), Jeffers (2019), Johnson, Lavetti, and Lipsitz (2020), for example.

[24]Note that I don't directly calculate skill-specific wage markdowns with the model. If broken down by skill, a model-implied dynamic markdown wedge would be much larger for most skilled workers, who are the most costly to replace. Hence wage markdown differences between high- and low-skill workers would likely not be quite as large as reduced-form empirical elasticity estimates seem to imply. But elasticity differences are still large enough that it's unlikely this would entirely reverse the effect of skilled workers' lower supply elasticities.

labor market power suggests such tradeoffs also exist in labor markets. For example, if antitrust actions that break up dominant firms in labor markets disincentivize investment that improves worker productivity or reduce the efficiency of their production technologies, the resulting productivity loss could destroy valuable worker-firm matches and reduce welfare, despite lowering aggregate monopsony power.

On the other hand, working to reduce artificial barriers to worker mobility targets the types of forces generating labor market power that are more clearly negative to workers from a welfare perspective. For example, antitrust actions which actively prevent employers from colluding to reduce worker mobility may be helpful. There is some precedent for such legal action already, as the US Department of Justice Antitrust Division filed a suit against a number of prominent Silicon Valley tech firms in 2010 for colluding from 2005 to 2009 to not poach each others' employees; a settlement forcing the firms to end this anti-competitive practice was quickly reached.[25] While reducing the scope of non-compete agreements and similar contracts restricting worker mobility may help raise wages, this may also generate countervailing effects of reducing investment rates among incumbent firms while increasing firm entry, and so the aggregate welfare implications of doing this are still unclear.[26]

5.3 Alternative 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.7.1 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).

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

[25]See court documents for the case "U.S. V. Adobe Systems, Inc., Et Al." found at https://www.justice.gov/atr/case/us-v-adobe-systems-inc-et-al.

[26]For example Jeffers (2019) shows that non-compete agreements can reduce corporate investment, and Ferres, Kankanalli, and Muthukrishnan (2023) find that the end of the no-poach Silicon Valley no-poach agreements negatively impacted colluding firms' innovation.

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, contributing to the decline in the aggregate labor share. The overall value of wage markdowns is substantial. The average firm earns about a third of its capital income from wage markdowns, and wage markdowns are worth about 40% of capital income in aggregate. Consistent with economic rents earned from labor market power, productive firms have higher valuation ratios but not significantly higher investment rates.

Although my findings suggest that labor market power generates substantial value to productive firms in the form of economic rents, my results suggest some caution in interpreting welfare or policy implications. In my model and in the data, productive firms actually pay higher wages despite having larger markdowns, since they operate on a steeper portion of the labor supply curve. Because of the opportunities they afford their employees–especially in the form of higher wages—highly productive firms are likely to be coveted employers.

Skilled workers play an important role in creating these patterns, as they face the lowest supply elasticities of any skill group. There is some evidence for an increased role of firm-specific human capital, but artificial barriers to mobility—such as the rising prevalence of non-compete agreements and labor market collusion between successful firms—may also have played a part. Exploring these possibilities in more depth could be a particularly useful avenue for further research.

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 D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics, Volume 4. Amsterdam: Elsevier-North*, pp. 1043–1171.

Acemoglu, D. and W. B. Hawkins (2014). Search with multi-worker firms. *Theoretical Economics 9*(3), 583–628.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2022). Good rents versus bad rents: R&d misallocation and growth.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-u relationship. *The Quarterly Journal of Economics 120*(2), 701–728.

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.

Autor, D. H. and D. Dorn (2013, August). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review 103*(5), 1553–97.

Azar, J., I. Marinescu, and M. Steinbaum (2022). Labor market concentration. *Journal of Human Resources 57*(S), S167–S199.

Balasubramanian, N., J. W. Chang, M. Sakakibara, J. Sivadasan, and E. Starr (2020). Locked in? the enforceability of covenants not to compete and the careers of high-tech workers. *Journal of Human Resources*.

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 (2020, August). Monopsony in movers: The elasticity of labor supply to firm wage policies. Working Paper 27755, National Bureau of Economic Research.

Belo, F., J. Li, X. Lin, and X. Zhao (2017, 07). Labor-Force Heterogeneity and Asset Prices: The Importance of Skilled Labor. *The Review of Financial Studies 30*(10), 3669–3709.

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 (2018, February). Strong employers and weak employees: How does employer concentration affect wages? Working Paper 24307, National Bureau of Economic Research.

Berger, D. W., K. F. Herkenhoff, and S. Mongey (2021, May). Labor market power. Working Paper 25719, National Bureau of Economic Research.

Boudoukh, J., R. Feldman, S. Kogan, and M. Richardson (2018, 07). Information, Trading, and Volatility: Evidence from Firm-Specific News. *The Review of Financial Studies 32*(3), 992–1033.

Brardsen, 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., 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., E. Mattana, S. Salgado, and M. Xu (2023). Wage setting and passthrough: The role of market power, production technology and adjustment costs.

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.

De Ridder, M. (2019). Market power and innovation in the intangible economy.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2020). Changing business dynamism and productivity: Shocks versus responsiveness. *The American Economic Review 110*(12), pp. 3952–3990.

Diamond, R. (2016, March). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review 106*(3), 479–524.

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. Papanikoloau (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.

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.

Garmaise, M. J. (2011). Ties that truly bind: Noncompetition agreements, executive compensation, and firm investment. *Journal of Law, Economics, and Organization 27*(2), 376–425.

Goolsbee, A. and C. Syverson (2023). Monopsony power in higher education: A tale of two tracks. *Journal of Labor Economics 41*(S1), S257–S290.

Gouin-Bonenfant, E. (2020). Productivity dispersion, between-firm competition, and the labor share.

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. Xioalan (2019). Capital share dynamics when firms insure workers. *The Journal of Finance 74*(4), 1707–1751.

Hazell, J., C. Patterson, H. Sarsons, and B. Taska (2022, November). National wage setting. Working Paper 30623, National Bureau of Economic Research.

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.

Jager, S. and J. Heining (2019). How substitutable are workers? evidence from worker deaths.

Jarosch, G., J. S. Nimczik, and I. Sorkin (2019, September). Granular search, market structure, and wages. Working Paper 26239, National Bureau of Economic Research.

Jeffers, J. (2019). The impact of restricting labor mobility on corporate investment and entrepreneurship.

Johnson, M. S., K. Lavetti, and M. Lipsitz (2020). The labor market effects of legal restrictions on worker mobility.

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 (2020, June). Imperfect competition and rents in labor and product markets: The case of the construction industry. Working Paper 27325, National Bureau of Economic Research.

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 (2019, June). Imperfect competition, compensating differentials and rent sharing in the u.s. labor market. Working Paper 25954, National Bureau of Economic Research.

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. (2023). Does monopsony matter for innovation?

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.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets.* Princeton University Press.

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. (2022). Labor market power, superstar firms, and between-firm wage (in) equality. *Superstar Firms, and Between-Firm Wage (In) Equality.*

Molloy, R., C. L. Smith, R. Trezzi, and A. Wozniak (2016). Understanding declining fluidity in the u.s. labor market. *Brookings Papers on Economic Activity*, 183–237.

Mueller, H. M., P. P. Ouimet, and E. Simintzi (2017, 05). Within-Firm Pay Inequality. *The Review of Financial Studies 30*(10), 3605–3635.

Neuhierl, A., A. Scherbina, and B. Schiusene (2013). Market reaction to corporate press releases. *The Journal of Financial and Quantitative Analysis 48*(4), 1207–1240.

Nimczik, J. (2020). Job mobility networks and data-driven labor markets.

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.

Prager, E. and M. Schmitt (2021, February). Employer consolidation and wages: Evidence from hospitals. *American Economic Review 111*(2), 397–427.

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. (2018). Labor market concentration, earnings inequality, and earnings mobility.

Robinson, J. (1969). *The Economics of Imperfect Competition.* Macmillan.

Rouwenhourst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In *Frontiers of Business Cycle Research*, pp. 294–330. Princeton Univesity Press.

Schubert, G., A. Stansbury, and B. Taska (2020). 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.*

Sokolova, A. and T. Sorensen (2021). Monopsony in labor markets: A meta-analysis. *ILR Review 74*(1), 27–55.

Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter (2018, 10). Firming Up Inequality. *The Quarterly Journal of Economics 134*(1), 1–50.

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.

Starr, E. P., J. Prescott, and N. D. Bishara (2021). Noncompete agreements in the us labor force. *The Journal of Law and Economics 64*(1), 53–84.

Stole, L. A. and J. Zwiebel (1996). Intra-firm bargaining under non-binding contracts. *The Review of Economic Studies 63*(3), 375–410.

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.

Vasicek, O. A. (1973). A note on using cross-sectional information in bayesian estimation of security betas. *The Journal of Finance 28*(5), 1233–1239.

Vuolteenaho, T. (2002). What drives firm-level stock returns? *The Journal of Finance 57*(1), 233–264.

Yeh, C., C. Macaluso, and B. Hershbein (2022, July). Monopsony in the us labor market. *American Economic Review 112*(7), 2099–2138.

İmrohoroğlu, A. and S. Tüzel (2014). Firm-level productivity, risk, and return. *Management Science 60*(8), 2073–2090.

Figures

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

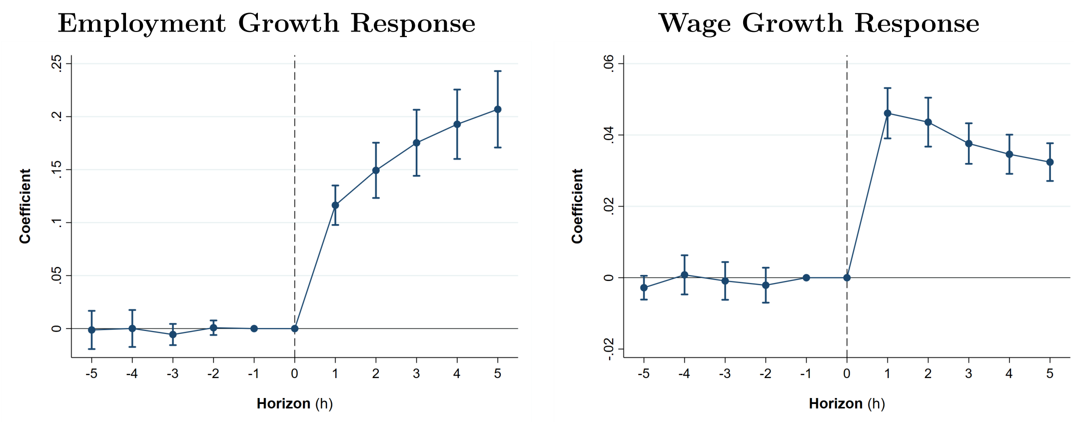

Note: This figures shows the employment and wage responses from estimating (6) in the main text for $h = -5$ to 5 years. The dependent variables are the growth average wages for the entire firm and firm Compustat employment. Confidence intervals are based off standard errors double clustered by industry and year.

Figure 2: Decline in Labor Supply Elasticities: Labor Productivity Sorts, 1991-2002 Versus 2003-2014

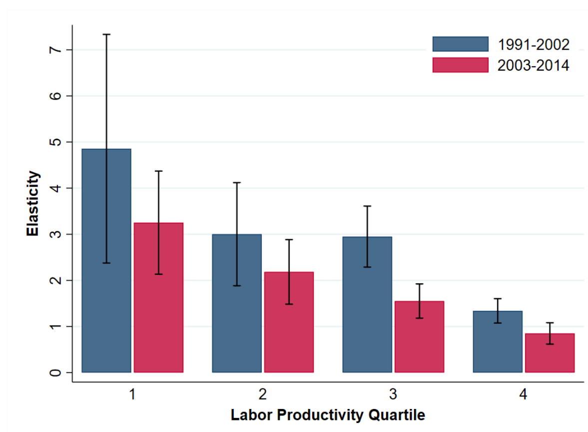

Note: This table reports estimates of estimates of the labor supply elasticity implied by estimating the employment and wage responses to stock returns, as in equation (7) of the main text, with estimates computed separately for the 1991-2002 and 2003-2014 subperiods. The main independent variable is the firm's own stock return in 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 intervaals are based off standard errors clustered by industry and year. See main text for further details.

Figure 3: The Downward Trend the Cross-Sectional Gap in Labor Shares and Spread in Estimated Markdowns Between High- and Low-Labor Productivity Firms Moves Together Over Time

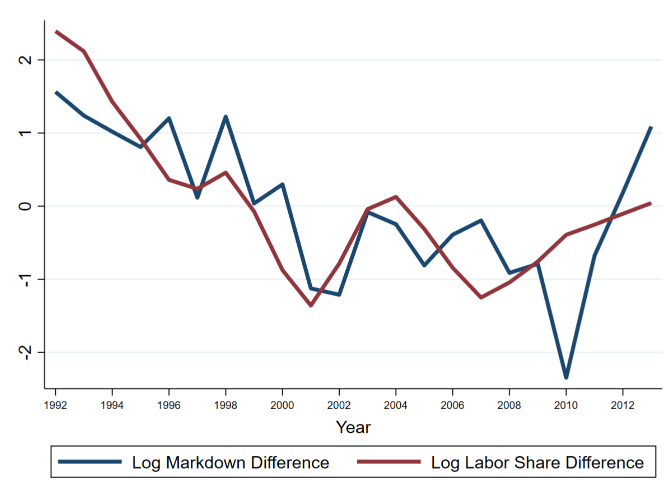

Note: This figure shows differences elasticity implied log markdowns between top- and bottom-quartile labor productivity firms from (8) in the main text, and log labor share differences following (9). The elasticity for a given year is for the moving window centered at that year (except for the start and end years, which are respectively the first or last year of 2-year window). 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.

Figure 4: Model Implied Markdowns

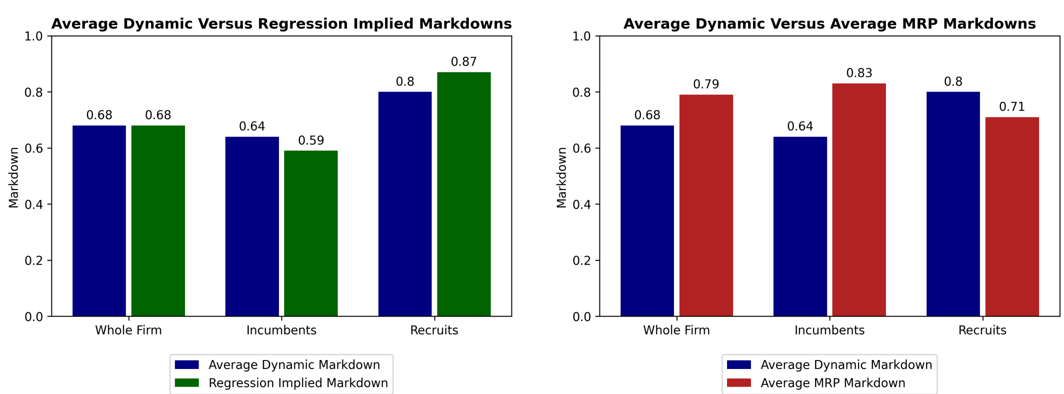

Note: This figure shows markdowns implied by the calibration of the model. The dynamic markdown is the markdown adjusted for model dynamics due to adjustment costs, and is given by (19) in the main text. The regression implied markdown is markdown implied by estimating the supply elasticity by running the regression of model employment/wage growth on stock returns and taking the ratio of the two coefficients. The MRP markdown is the markdown from marginal revenue product. See section 3 in text for details.

Tables

Table 1: Labor Supply and Rent-Elasticities Implied by Employment, Wage, and Value-Added Per Worker Responses Stock Returns

By Worker Type			By Labor Productivity Quartile			
Panel A: Employment Growth Response to Stock Return						
All Workers	Incumbents	Recruits	Q1	Q2	Q3	Q4
0.12	0.091	0.20	0.12	0.089	0.12	0.11
(0.0095)	(0.0090)	(0.016)	(0.013)	(0.0097)	(0.011)	(0.011)
Panel B: Wage Growth Response to Stock Return						
All Workers	Incumbents	Recruits	Q1	Q2	Q3	Q4
0.046	0.053	0.028	0.027	0.033	0.051	0.091
(0.0036)	(0.0037)	(0.0021)	(0.0032)	(0.0029)	(0.0040)	(0.0097)
Panel C: Implied Labor Supply Elasticity						
All Workers	Incumbents	Recruits	Q1	Q2	Q3	Q4
2.52	1.73	7.19	4.32	2.71	2.34	1.17
(0.30)	(0.22)	(0.49)	(0.87)	(0.43)	(0.28)	(0.15)
Panel D: VA/Worker Growth Response to Stock Return						
All Workers	Incumbents	Recruits	Q1	Q2	Q3	Q4
0.18	0.18	0.18	0.17	0.15	0.17	0.26
(0.017)	(0.017)	(0.017)	(0.021)	(0.016)	(0.026)	(0.017)
Panel E: Implied Rent-Sharing Passthrough Elasticity						
All Workers	Incumbents	Recruits	Q1	Q2	Q3	Q4
0.25	0.29	0.15	0.16	0.22	0.30	0.35

Note: This table reports estimates of equation (5) for all workers, incumbent, workers, and new recruits in the first three columns, and shows estimates of (7) in the last four columns. The independent variable is the firm's own stock return in 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 in the final four columns. In panel A the dependent variable is the one-year employment growth rate; in panel B it is the one-year growth rate in the firm's average wage; in panel D the dependent variable is the growth rate in firm value-added per worker, and the lagged employment and growth controls are replaced by a control for lagged value-added per worker growth. The labor supply elasticities in panel C are given by the ratio between estimates in panel A to the estimates in panel B; the rent-sharing elasticities in panel E are given by the ratio between the estimates in panel B to the estimates in panel D. Standard errors in parentheses are clustered by industry and year, and elasticity standard errors in panel C are computed via two-stage least squares (I do not report standard errors in panel E because the controls in panels B and D are slightly different). The p-value on the difference between the top and bottom labor productivity quartile supply elasticity estimates in panel C is 0.001. The sample size for the first three columns is 45,500 and 43,500 in the final four columns (both rounded per Census disclosure rules). See main text for further details.

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	log Q (Tot)	log M/B	ROA_t	ROA_{t+1}
log VA/Worker	-0.56	0.52	0.25	0.11	0.10
	(0.04)	(0.05)	(0.03)	(0.01)	(0.01)
Size Controls	X	X	X	X	X
Industry X Year FE	X	X	X	X	X
N	57500	48500	55000	57500	53500
R^2 (within)	0.48	0.13	0.07	0.17	0.13

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 (income before extraordinary items plus deprecation over total assets). Operating performance and valuation ratios are winsorized at the 1% level by year. See section 2.3 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: Model Calibration and Model-implied Moments versus Data

Panel A: Model Calibration

Parameter	Explanation	Value
Externally Calibrated		
L	Labor force size	1.0
r	Discount rate	0.02
$\overline{w}$	Wage upper bound	1.0
ρ_z	Productivity persistence	0.9
$\bar{z}$	Fixed labor productivity	1.0
σ_c	Adjustment cost volatility	1.0
Internally Calibrated		
α	1 - Labor returns to scale	0.22
β	Supply elasticity shifter	0.38
b	Reservation wage	0.545
σ_z	Productivity volatility	0.145
$\bar{C}$	Adjustment cost intensity	0.57

Panel B: Targeted Moments

Moment Name	Model	Data
Separations Rate	0.31	0.30
Incumbent Premium	0.15	0.15
Supply Elasticity	2.09	2.52
Elasticity (Inc.)	1.44	1.73
Elasticity (Rec.)	6.66	7.19
Inc/Rec Wage Pass. Ratio	1.93	1.86
Log Labor Share	-0.50	-0.51

Panel C: Heterogeneity by Productivity

Moment Name	Model	Data
Supply Elasticity (Q1)	4.85	4.32
Supply Elasticity (Q2)	2.53	2.71
Supply Elasticity (Q3)	1.80	2.34
Supply Elasticity (Q4)	1.36	1.17
Log Labor Share (Q1)	-0.31	-0.22
Log Labor Share (Q2)	-0.46	-0.40
Log Labor Share (Q3)	-0.56	-0.54
Log Labor Share (Q4)	-0.68	-0.94

Panel D: Rent-Sharing Elasticities

Elasticity for:	Model	Data
All workers	0.30	0.25
Incumbents	0.36	0.29
Recruits	0.19	0.15
Productivity (Q1)	0.16	0.16
Productivity (Q2)	0.22	0.22
Productivity (Q3)	0.34	0.30
Productivity (Q4)	0.38	0.35

Note: Panel A of this table shows the parameter calibration for the model in section 3 of the text. Externally calibrated parameters are calibrated beforehand without trying to match target model moments, while internally calibrated parameters are explicitly set in order to match target moments. Panels B through D of this table compare model implied moments with their empirical counterparts; panel B reports targeted moments, C reports heterogeneous supply elasticities and average log labor shares by labor productivity quartile; and panel D computes rent-sharing elasticities in the model. “Inc" and “Rec" denote incumbents and recruits, respectively. In order to mimic the empirical estimates, the model supply elasticities are computed by running a regression of model-generated stock returns on employment and wage growth and taking the ratio of the model coefficients, and rent-sharing elasticities are estimated by taking the ratio of the wage growth to log value-added per worker growth coefficients on stock returns in the model. See section 3 in text for details.

Table 4: Adjusted Markdowns for Firms Sorted on Labor Productivity, as Implied By Empirical Elasticity Estimates and Markdown Wedge ν From Model

	Elasticity	Unadjusted Markdown	Adjusted Markdown
Overall	2.52	0.72	0.83
Quartile 1	4.32	0.81	0.94
Quartile 2	2.71	0.73	0.85
Quartile 3	2.34	0.70	0.81
Quartile 4	1.17	0.54	0.62

Note: This table compares markdowns implied by my empirical supply elasticity estimates before adjusting for the model markdown wedge ν estimated from the model and after. For supply elasticity ϵ, the unadjusted markdown is $\epsilon/(1+\epsilon)$; the adjusted markdown is $\nu \times \epsilon/(1+\epsilon)$.

Table 5: Valuing Labor Market Power—Wage Markdowns as a Fraction of Operating Income

Panel A: Full Sample (1991-2014)

	Overall	Low Productivity	High Productivity
Mean Firm-Level Operating Income Share	0.34	0.17	0.43
Median Firm-Level Operating Income Share	0.30	0.08	0.49
Mean Aggregate Operating Income Share	0.40	0.10	0.45

Panel B: By Subperiod (1991-2002 and 2003-2014)

	Overall	Low Productivity	High Productivity
1991-2002			
Mean Firm-Level Operating Income Share	0.34	0.19	0.42
Median Firm-Level Operating Income Share	0.29	0.10	0.45
Mean Aggregate Operating Income Share	0.39	0.12	0.43
2003-2014			
Mean Firm-Level Operating Income Share	0.35	0.14	0.44
Median Firm-Level Operating Income Share	0.32	0.07	0.52
Mean Aggregate Operating Income Share	0.41	0.07	0.46

Note: This table shows the dollar value of wage markdowns as a fraction of operating income; these are broken down by the full sample period in panel A and by subperiod in panel B. See section 4.1 of the main text for more details.

Table 6: Differences in Labor Shares Explained by Markdowns, Cross-Section and Time Series

Panel A: High Minus Low Productivity Cross-Sectional Average Labor Share Spread

	Low Productivity	High Productivity	High - Low	% Due to ΔMarkdown
Log Markdown	-0.05	-0.46	-0.41	
Avg Log LShare ($\widehat{\textsc{xlr}}$)	-0.18	-0.83	-0.65	63.5
Avg Log LShare ($\textsc{labex}$)	-0.22	-0.94	-0.72	57.1

Panel B: Time Series Change in Aggregate Labor Share

	$1991-2002$	$2003-2014$	$\widetilde{2003-2014}$	$\Delta\log(\text{Lshare})$	$\Delta\widetilde{\log(\text{Lshare})}$	% Due to ΔMarkdown
Agg LShare ($\widehat{\textsc{xlr}}$)	0.579	0.524	0.554	-0.10	-0.044	55.8
Agg LShare ($\textsc{labex}$)	0.506	0.485	0.515	-0.042	0.018	143.2

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 5. 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 column $\widetilde{2003-2014}$ denotes the counterfactual average labor share for the 2003-2014 period if markdowns were held constant at their 1991-2002 estimates, as defined in (31) in the main text. The column $\Delta\log(\text{Lshare})$ gives the actual log change in the average labor share between the two periods, while $\Delta\widetilde{\log(\text{Lshare})}$ gives the counterfactual log labor share change holding markdowns constant. The last column of panel B gives the fraction of the observed change in the labor share that is attributable to wage markdown changes. See section 4.2 in main text for further details.

Table 7: Fraction of Variation in Average Valuation Ratios across Labor Productivity Quartiles Explained by Wage Markdowns

	Avg Q by Productivity Quartile				**% Explained by Counterfactual**			
	1	2	3	4	4-1 Diff	%	σ^2(Avg Q)	%
Counterfactual	0.49	0.52	0.71	0.83	0.34		0.03	
Predicted	0.64	0.89	1.12	1.58	0.94	63.9	0.16	83.9
Actual	0.73	0.79	1.04	1.70	0.97	65.2	0.20	87.0

Note: The first row of this table shows the average counterfactual Peters and Taylor (2017) total Tobin's Q valuation ratios based off firm operating income-to-capital ratios by labor productivity quartile. Counterfactual valuation ratios replace actual operating income-to-capital ratios with counterfactual operating income after subtracting off the dollar value of wage markdowns, as in equation (A.13). The next two rows of this table respectively show the equation (A.12) regression-predicted average and the actual average Q ratios for the given productivity group. The last 4 columns of the table examine what percentage of either the top - bottom productivity quartile gap or the variance of average valuations across productivity quartiles can be explained by replacing actual valuations with counterfactual predicted valuations, calculated as $100 \times (\text{Baseline} - \text{Counterfactual})/\text{Baseline}$ for Baseline = Predicted, Actual. The valuation ratios, actual operating income-to-capital ratios, and counterfactual actual operating income-to-capital ratios are all cross-sectionally winsorized at the 1% level.

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.

$$\mathrm{VA}_{j,t} = \mathrm{OIBDP}_{j,t} + \Delta\mathrm{INVFG}_{j,t} + \mathrm{LABEX}_{j,t} \tag{A.1}$$

Changes in inventories are set to zero when missing. Here

$$\mathrm{LABEX}_{j,t} = \widetilde{W}_{j,t} \times (\mathrm{EMP}_{j,t} + \mathrm{EMP}_{i,t-1})/2 \tag{A.2}$$

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

$$\log(\mathrm{VA}/\mathrm{Worker})_{j,t} = \log(\mathrm{VA}_{j,t}/\mathrm{EMP}_{j,t}) \tag{A.3}$$

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

$$\mathrm{LSHARE}_{j,t} = \frac{\mathrm{LABEX}_{j,t}}{\mathrm{VA}_{j,t}} \tag{A.4}$$

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 et al. (2013). I then estimate the wage decomposition

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

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 $N^{ins}_{j,t}$ denote the total number of workers mapped to firm j that are in the sample of (A.5) in year t. Let $N^{not-ins}_{j,t}$ denote the number of unique individuals that show up on the payrolss 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

$$\widetilde{W}_{j,t} = \frac{\sum_i \widetilde{w}_{ijt}}{N^{ins}_{j,t} + N^{not-ins}_{j,t}} \tag{A.6}$$

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:

$$W_{j,t} = \frac{\sum_{i \in \Gamma_{j,t}} w_{ijt}}{N^{ins}_{j,t}} \tag{A.7}$$

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 2. 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:

$$s_{j,t} = \log\left(\frac{\sum_{i\in\Gamma_{j,t}} \exp(\widehat{\alpha}_{i,t})}{N^{ins}_{j,t}}\right) \tag{A.8}$$

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 and 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

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

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

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

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

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

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:

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

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:

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

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 $\widehat{XLR}_{j,t} = \exp\left(\widehat{\alpha} + \widehat{\alpha}_{I(j),t} + \widehat{\beta_t}\log(LABEX_{j,t})\right)$. 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 a much closer to actual $XLR_{j,t}$. Because of this I use the exponential of the log to create a strictly positive predicted $XLR_{j,t}$. In-sample $\widehat{XLR}_{j,t}$ has a correlation of 0.95 with actual $XLR_{j,t}$.

A.4 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:

$$Q = \frac{1}{r-g}\frac{\Pi}{K} - \frac{1}{r-g}R + q \tag{A.11}$$

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 (A.11) 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:

$$Q_{j,t} = \beta_t \times \left(\frac{OI_{j,t}}{K_{j,t}} \right) + \alpha_{I(j),t} + \epsilon_{j,t} \tag{A.12}$$

I estimate (A.12) period-by-period. The industry fixed effects $\alpha_{I(j),t}$ impose that the user cost of capital and marginal q are constant within the industry-year. I then take the $\widehat{\beta_t}$ 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 4:

$$\widetilde{Q}_{j,t} = \widehat{\beta_t} \times \left(\frac{\widetilde{OI}_{j,t}}{K_{j,t}} \right) + \widehat{\alpha}_{I(j),t} \tag{A.13}$$

where $\widetilde{OI}_{j,t} = OI_{j,t} - \left(\widetilde{LABEX}_{j,t} - LABEX_{j,t} \right)$, with $\widetilde{LABEX}_{j,t}$ as defined in (25).

In Table 7 I use (A.13) to predict counterfactual average Peters and Taylor (2017) average total Tobin's Q valuation ratios by labor productivity quartile. I also report the averages as predicted by operating profitability in (A.11) along with the actual data averages. The last four columns of the table report the top-minus-bottom productivity quartile gap in average valuations and the variance of average valuations across the four productivity groups, as well as the percentage that can be explained by the counterfactual valuation ratios. This percentage is calculated as $100 \times (\text{Baseline} - \text{Counterfactual})/\text{Baseline}$ for $\text{Baseline} = \text{Predicted}, \text{Actual}$. Table 7 shows that the majority of variation in average valuation ratios across labor productivity quartiles (either as measured by the cross-sectional spread or variance of average valuations across productivity groups) can be explained by removing the capital income coming from wage markdowns.

8 Internet Appendix

IA.1 Elasticity Estimate Bias in a Simple Model of the Labor Market

Similar to Card et al. (2018) or Lamadon et al. (2019), suppose that L workers choose employers among a market of N firms, and normalize $L = 1$ for simplicity. Worker i's utility from working at firm j is increasing in firm-specific amenities $a_{j,t}$, the log of the wage offer $w_{j,t}$, and an unobservable taste shock $\epsilon_{i,j,t}$.

$$u_{i,j,t} = \varepsilon \log(w_{j,t}) + a_{j,t} + \epsilon_{i,j,t} \tag{IA.1}$$

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

$$L(w_{j,t}, a_{j,t}) = \lambda_t^{-1} \exp(a_{j,t}) w_{j,t}^{\varepsilon} \tag{IA.2}$$

Like Card et al. (2018) and Lamadon et al. (2019), 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:

$$V_{j,t} = \max_{w_{j,t}} Z_{j,t}(L_{j,t})^{1-\alpha} - w_{j,t} L_{j,t} \tag{IA.3}$$

subject to the functional form of the labor supply curve (IA.2). Denote the wage markdown by $\mu \equiv \frac{\varepsilon}{\varepsilon+1}$ and define constant $c \equiv \frac{1}{1+\varepsilon\alpha} \log\left(\mu(1-\alpha)\right)$. Solving (IA.3) yields the expressions for the log optimal employment, wage, and firm value:

$$\log(L_{j,t}) = \varepsilon c + \frac{\varepsilon}{1+\varepsilon\alpha} \log(Z_{j,t}) + \frac{1}{1+\varepsilon\alpha} \log(A_{j,t}) \tag{IA.4}$$

$$\log(w_{j,t}) = c + \frac{1}{1+\varepsilon\alpha} \log(Z_{j,t}) - \frac{\alpha}{1+\varepsilon\alpha} \log(A_{j,t}) \tag{IA.5}$$

$$\log(V_{j,t}) = \log(1-(1-\alpha)\mu) + (1-\alpha)\varepsilon c + \frac{1+\varepsilon}{(1+\varepsilon\alpha)} \log(Z_{j,t}) + \frac{1-\alpha}{1+\varepsilon\alpha} \log(A_{j,t}) \tag{IA.6}$$

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

$$\hat{\varepsilon}_{IV}^{\log V} = \frac{\varepsilon(1+\varepsilon)\sigma_z^2 + (1-\alpha)\sigma_a^2 + \varepsilon(1-\alpha)\sigma_{az} + (1+\varepsilon)\sigma_{az}}{(1+\varepsilon)\sigma_z^2 + (1-\alpha)\sigma_{az} - \alpha(1+\varepsilon)\sigma_{az} - \alpha(1-\alpha)\sigma_a^2} \tag{IA.7}$$

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.7) collapses to the true supply elasticity ε.

For (IA.7) to represent an upper bound on the supply elasticity—which implies a conservative estimate of the magnitude of wage markdowns—we must have

$$\rho_{az}\sigma_z > -\frac{(1-\alpha)}{1+\varepsilon}\sigma_a \tag{IA.8}$$

where ρ_{az} is the correlation of $\log(A_{j,t})$ and $\log(Z_{j,t})$. The intuition behind (IA.8) 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.8) is market-specific productivity shocks. This is because $A_{j,t}$ is decreasing in the market-wide wage index $\lambda_t = \left(\sum_{j'=1}^{N} \exp(a_{j',t}) w_{j',t}^{\varepsilon}\right)$:

$$A_{j,t} = \lambda_t^{-1} \exp(a_{j,t}) = \left(\sum_{j'=1}^{N} \exp(a_{j',t}) w_{j',t}^{\varepsilon}\right)^{-1} \exp(a_{j,t}) \tag{IA.9}$$

Suppose that firm-specific productivity $Z_{j,t}$ is given by an aggregate market component and an idiosyncratic component: $Z_{j,t} = \widetilde{X}_t X_{j,t}$. Equation (IA.5) then implies that the wage offer can be expressed as

$$w_{j,t} = C_{j,t} \widetilde{X}_t^{\frac{1}{1+\alpha\varepsilon}} \lambda_t^{\frac{\alpha}{1+\alpha\varepsilon}} \tag{IA.10}$$

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

$$\lambda_t^{\frac{1}{1+\alpha\varepsilon}} = \left(\sum_{j'} \widetilde{C}_{j',t}\right) \widetilde{X}_t^{\frac{\varepsilon}{1+\alpha\varepsilon}} \equiv \widetilde{C}_t \widetilde{X}_t^{\frac{\varepsilon}{1+\alpha\varepsilon}} \tag{IA.11}$$

or

$$\lambda_t = \widetilde{C}_t^{1+\alpha\varepsilon} \widetilde{X}_t^{\varepsilon} \tag{IA.12}$$

Hence $\log(A_{j,t}) = a_{j,t} - (1 + \alpha\varepsilon) \log\left(\widetilde{C}_t\right) - \varepsilon \log\left(\widetilde{X}_t\right)$ is decreasing in $\log(\widetilde{X}_t)$. If enough of the variance in $\log(A_{j,t})$ is driven by $\log(\widetilde{X}_t)$ this can reverse the inequality in (IA.8), 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 $\log(\widetilde{C}_t)$ and $\log(C_{j,t})$ is approximately zero and so I don't consider the impact of this term. The $\log(\widetilde{C}_t)$ term nets out with market-wide fixed effects along with $\log(\widetilde{X}_t)$. 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.8) becomes highly plausible. When market shocks are netted out, (IA.8) 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.8) holds, implying my estimates would be conservative if biased at all.

Now, compare the IV estimate of the elasticity obtained in equation (IA.7) with the OLS estimate from regressing log employment on log wages:

$$\hat{\varepsilon}_{OLS} = \frac{\varepsilon\sigma_z^2 - \alpha\sigma_a^2 + (1 - \varepsilon\alpha)\sigma_{az}}{\sigma_z^2 + \alpha^2\sigma_a^2 - 2\alpha\sigma_{az}} \tag{IA.13}$$

In contrast to the IV estimate, the OLS elasticity estimate in (IA.13) can be biased up or down depending on parameters, even after invoking assumption (IA.8). 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.

IA.2 Wage Markdowns and Rent Sharing Under a Multilateral Bargaining Framework

IA.2.1 Wage Markdowns under Multilateral Bargaining versus Monopsonistic Wage Posting

I 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:

$$\max_{L} F(L) - c(L) - W(L)L \tag{IA.14}$$

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

$$W(L) = (1-\eta)b + \gamma\eta F_L(L) \tag{IA.15}$$

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:

$$F_L(L) = c_L(L) + W_L(L)L + W(L) \tag{IA.16}$$

I solve for the expression for the quasi labor supply elasticity

$$\frac{d\log L^*}{d\log W^*} \tag{IA.17}$$

where L^* is the solution to the firm's problem (IA.14) and $W^* = W(L^*)$. I do so by first solving for the inverse labor supply elasticity through a shock to Z:

$$\frac{dW^*}{dL^*}\frac{L^*}{W^*} = \left(\frac{\partial W^*}{\partial Z}\frac{\partial Z}{\partial L^*} + \frac{\partial W^*}{\partial L^*}\right)\frac{L^*}{W^*} \tag{IA.18}$$

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.14). The components of (IA.16) are given by

$$F_L = Z(1-\alpha)L^{-\alpha} \tag{IA.19}$$

$$c_L = kL^{\theta} \tag{IA.20}$$

$$W(L) = (1-\eta)b + \gamma\eta Z(1-\alpha)L^{-\alpha} \tag{IA.21}$$

$$W_L L = -\alpha\gamma\eta Z(1-\alpha)L^{-\alpha} \tag{IA.22}$$

Collecting terms for (IA.16):

$$(1-\alpha)ZL^{-\alpha} = kL^{\theta} - \alpha\gamma\eta(1-\alpha)ZL^{-\alpha} + (1-\eta)b + \gamma\eta(1-\alpha)ZL^{-\alpha} \tag{IA.23}$$

Applying the implicit function theorem to (IA.23) yields the expression for $\frac{\partial Z}{\partial L}$

$$\frac{\partial Z}{\partial L} = \frac{\theta kL^{\theta+\alpha} + \alpha(1-\alpha)Z - \alpha\eta\gamma(1-\alpha)Z + \alpha^2(1-\alpha)\eta\gamma Z}{\left[(1-\alpha) - \gamma\eta(1-\alpha) + \alpha\eta\gamma(1-\alpha)\right]L} \tag{IA.24}$$

Next, differentiating (IA.21) with respect to Z:

$$\frac{\partial W}{\partial Z} = \eta\gamma(1-\alpha)L^{-\alpha} \tag{IA.25}$$

Finally, differentiating (IA.21) with respect to L gives

$$\frac{\partial W}{\partial L} = -\alpha\eta\gamma(1-\alpha)L^{-\alpha-1} \tag{IA.26}$$

Plugging in (IA.24), (IA.25), and (IA.26) to (IA.18), applying the definition of γ, and algebraic manipulation results in the inverse quasi-elasticity:

$$\frac{d\log W}{d\log L} = \frac{\theta kL^{\theta}}{\frac{1-\eta}{\eta}W} \tag{IA.27}$$

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.27) yields the desired quasi-elasticity of labor supply

$$\frac{d\log L}{d\log W} = \frac{W}{\theta(W-b)} \equiv \tilde{\varepsilon} \tag{IA.28}$$

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. (2013), Kline et al. (2019), and in my framework in section 3 of the main text:

$$\hat{\varepsilon}_{Wage\ posting} \equiv \frac{\beta W}{(W-b)} \tag{IA.29}$$

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:

$$\text{Quasi Markdown} = \frac{\tilde{\varepsilon}}{1+\tilde{\varepsilon}} = \frac{W}{(1-(1+\theta)\eta)b + (1+\theta)\eta\gamma F_L(L)} \tag{IA.30}$$

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

$$\text{True Markdown} = \frac{W}{F_L(L)} \tag{IA.31}$$

The true markdown (IA.31) under multilateral bargaining also implies that wage markdowns are increasing in labor productivity, similar to my model in section 3. The relative bias ratio in wage markdowns from assuming wage posting monopsony instead of multilateral bargaining is

$$\frac{\text{True Markdown}}{\text{Quasi Markdown}} = \frac{(1-(1+\theta)\eta)b + (1+\theta)\eta\gamma F_L(L)}{F_L(L)} \tag{IA.32}$$

To understand the size of the bias, consider the special case $1+\theta = 1/\eta$. Then (IA.32) 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.23) can be rearranged as follows:

$$\gamma F_L(L) = \frac{kL^\theta}{1-\eta} + b \tag{IA.33}$$

which implies $\gamma F_L(L) > b$.[27] This further implies that the numerator in (IA.32) is strictly increasing in $(1+\theta)$. Hence the ratio must be strictly decreasing as convexity θ decreases, and in turn smaller than γ when $1+\theta < 1/\eta$, and the ratio (IA.32) 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.2.2 Rent-Sharing Elasticities Under Multilateral Bargaining versus Wage Posting

The wage equation (IA.21) from the multilateral bargaining framework delivers the following wage elasticity with respect to output per worker:

$$\frac{d\log W}{d\log F/L} = \frac{\gamma\eta(1-\alpha)F/L}{(1-\eta)b + \gamma\eta(1-\alpha)F/L} \equiv \xi_{Bargaining} \tag{IA.34}$$

Again the connection with a wage posting model where labor supply is proportional to $(W-b)^\beta$ is readily apparent:

$$\xi_{Wageposting} = \frac{\frac{\beta}{1+\beta}(1-\alpha)F/L}{\frac{1}{1+\beta}b + \frac{\beta}{1+\beta}(1-\alpha)F/L} \tag{IA.35}$$

Both models deliver output-per-worker elasticities that are increasing in F/L but strictly less than 1.

I map F/L to the firm's value-added per worker in the data, where value-added is defined as the sum of operating income before depreciation plus changes in inventories and labor expenses, following common practice for Compustat firms (Donangelo et al., 2019). This empirical measure of value-added is comparable to the notion of gross surplus used by Kline et al. (2019), who take the sum of EBITDA (the Compustat variables for EBITDA and OIBDP are exactly the same) and labor expenses; following their assumption that firms do not claim deductions on hiring costs from their operating income, the elasticities defined above map into the wage/value-added per worker rent-sharing elasticities that I estimate

[27] Note that this rearrangment 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.

empirically. These equations also make clear that rent-sharing elasticities should *increase* in labor productivity rank, while labor supply elasticities decrease. The empirical estimates in the main text confirm both these predictions.

IA.3 Supply Elasticity Robustness Checks: Empirical Specifications and Data

IA.3.1 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.7 I present the elasticity estimates implied by estimating my baseline specification (5), as well as for a set of 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.7 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.7 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.7 I drop the industry-by-year fixed effects. Consistent with the discussion above and the simple model in internet appendix IA.1, 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.7 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.5 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.[28] See appendix Table IA.13 for more details on the comparative performance of different definitions of labor market boundaries. In column 3 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, 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 2.4 and 2.6, 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 consider my simple industry-by-year fixed effect specification to be a reasonable parsimonious baseline specification. 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.3.2 Alternative Labor Demand Shocks

I now perform a number of robustness checks in order to gauge how much unobserved firm-specific labor supply shocks that are correlated with stock returns may be biasing my estimates. I first estimate supply elasticities using several different proxies for firm labor demand shifters. Results are found in appendix Table IA.8. I report the estimated supply elasticity using the alternative measure, the supply elasticity from my baseline specification, and the p-value on the test that my baseline supply elasticity estimates are different from the given estimate. Since one the measures is available only for a subsample of the data, I re-estimate my baseline specification for the same subsample of the data before comparing elasticity estimates.

In the first column of IA.8 I use stock returns of firms' customers rather than the firms themselves as shock to labor demand. To test if firm-specific exposure to common stock market risk factors matters, in the next column I use return residuals backed out from a regression on the 5 Fama and French (2015) factors, plus the momentum risk factor. In the third column I use patent-induced shocks from Kogan, Papanikolaou, Seru, and Stoffman

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

(2017) as a shifter of firm labor demand; and finally, in the last column I construct firms announcement surprises following Daniel, Hirshleifer, and Sun (2019) and Chan, Jegadeesh, and Lakonishok (1996) in 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.

Each different labor demand shifter in Table IA.8 yields elasticity estimates that are very close quantitatively and statistically indistinguishable from my baseline estimates. Because each of these measures can be expected to covary differently with possible unobserved firm-specific labor supply shocks, this suggests that any bias from such shocks is not likely to have large a quantitative effect on my stock-return based supply elasticity estimates.

In the first column of Table IA.8 I show the response to stock returns of firms' customers using cleaned Compustat customer-supplier data provided by Wharton Research Data Services. 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. Still, a potential concern is that customers may also be labor market competitors, and so in unreported results I consider a version where I exclude customer-supplier links between firms who have hired from one another within a 5 year window centered at the current year, or who are in the same 3-digit NAICS industry. In both cases the estimated supply elasticity is nearly the exact same as the estimate based on own-firm stock return and is statistically indistinguishable, with both elasticity estimates for the customer return sample being close to 2.

In the second column of Table IA.8 I use the firms' cumulative abnormal log excess returns with respect to the Fama and French (2015) five factor model augmented with the momentum factor.[29] I construct stock return residuals as follows. In each year and for each firm i I regress daily log excess returns on the 6 risk factors:

$$\log(R_{i,t}) - \log(R_{f,t}) = \alpha_i^\tau + \beta_i^\tau F_t + \epsilon_{i,t} \tag{IA.36}$$

Here τ denotes the year corresponding to day t and F_t is a vector of risk factors. I take the estimated $\hat{\beta}_i$ for year τ and take these as the risk exposures for year $\tau + 1$ to get cumulative abnormal log returns. A literature starting with Vasicek (1973) suggests that out-of-sample

[29]All risk factor and risk-free rate data are obtained from Ken French's website.

risk exposures are more accurately estimated when applying Bayesian shrinkage to deflate the estimates towards a common average. I shrink the $\hat{\beta}_i$ estimates towards the cross-sectional average $E_\tau[\hat{\beta}_i]$, using the cross-sectional variance of the $\hat{\beta}_i$ estimates, $\text{Var}_\tau(\hat{\beta}_i^\tau)$, and the standard error of the own firm's estimate, $SE^2_{\beta_i}$, to get the shrinkage weights:

$$\widetilde{\beta_i^\tau} = (1-\omega)\hat{\beta}_i + \omega E_\tau[\hat{\beta}_i^\tau] \tag{IA.37}$$

where

$$\omega = \frac{SE^2_{\beta_i}}{SE^2_{\beta_i^\tau} + \text{Var}_\tau(\hat{\beta}_i^\tau)} \tag{IA.38}$$

I similarly deflate the intercepts α_i. Finally, the cumulative abnormal returns for year $\tau+1$ are given by

$$\text{Abnormal Return}_{i,\tau} = \sum_{t\in\tau} \log(R_{i,t}) - \log(R_{f,t}) - \widetilde{\beta_i^{\tau-1}} F_t - \widetilde{\alpha_i^{\tau-1}} \tag{IA.39}$$

I use Abnormal $\text{Return}_{i,\tau}$ in place of the excess return and re-estimate (5). I follow the same timimg convention as in mys baseline spec so that abnormal returns are aggregated from July until June of the following year. My inclusion of this measure relates primarily to the need to control for market shocks discussed previously and demonstrated in Table IA.7. In particular, a long literature in asset pricing argues that these factors represent systematic shocks, exposure to which demands compensation with higher returns. My inclusion of industry-by-year fixed effects implicitly imposes constant betas within industry on these systematic shocks, but it's possible the heterogeneity in exposure to systematic factors could matter. This procedure accordingly allows firms to have differential exposures to common factors. Using this measure again yields a very similar elasticity estimate to my baseline (2.34 relative to a baseline of 2.52, difference statistically insignificant).

In the third column I show the supply elasticity implied by employment and wage responses to patent induced shocks to firm value from Kogan et al. (2017), who show that the 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 predicted Kogan et al. (2017) patent values. Again I find an elasticity estimate that is close to my baseline.

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 idea here is to isolate 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 isolate price movements related to information about firm labor demand, and less likely to be related to firm-specific labor supply shocks, which would tend to be much more slow-moving. Though the measures are different, this follows a similar intuition to 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.

IA.3.3 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.14, 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).[30] 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.14. 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

[30] Thanks to Matthew Knepper for generously sharing his data.

my estimates. Meanwhile, Jeffers (2019) argues that firms use non-competes 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.14 I re-estimate supply elasticities based off (5) 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.14 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 columnn. 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.14 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.4 Robustness Checks for Labor Productivity-Sorted Supply Elasticity Estimates

IA.4.1 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 IA.3 could be reliance on unconditional labor productivity sorts that use between rather than within-industry variation. In Table IA.9 I instead sort firms into productivity quartiles within their 2-digit NAICS industry and re-estimate (7) for employment and wages to back out supply elasticities for workers of all skill levels. All my basic findings from Table IA.3 are unchanged for the within-industry sorts in IA.9, and all the elasticity point estimates are very similar.

In appendix Table IA.10 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 (7) 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:

$$\log(Y_{j,t+3}) - \log(Y_{j,t}) = \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} + \sum_{q=1}^{4} \mathbf{1}(q(i,t) = q) \times \beta_q \text{Stock Ret}_{j,t\to t+3} + \Gamma X_{j,t} + \epsilon_{j,t} \tag{IA.40}$$

Panel A of Table IA.10 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.10 I re-estimate (7) 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 elasticites. These elasticity estimates are not surprisingly a little higher than in Table IA.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.4.2 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.10 I sort on estimates of firm total factor productivity from İmrohoroğlu 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.10 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.[31] 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 LBD-based supply elasticities would increase the quantitative magnitude of my findings, if anything.

IA.4.2 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 IA.3 already alleviates these concerns in part, because productive firms face significantly lower supply elasticities even for

[31] LBD 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.

the least skilled workers, whose compensation is far less tied to equity-based incentive pay.[32]

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.11 I allow wage responses to be asymmetric for positive and negative excess returns for firms sorted on labor productivity by estimating the following:

$$\begin{aligned} \log(w_{j,t+1}) - \log(w_{j,t}) &= \alpha + \alpha_{q(j,t)} + \alpha_{I(j),t} \\ + \sum_{q=1}^{4} \mathbf{1}(q(i,t) = q) \times \Big[&\beta_{1,q} \text{Stock Ret}_{j,t\to t+1} \times \mathbf{1}\left(\text{Stock Ret}_{j,t\to t+1} \geq 0\right) \\ &+ \beta_{2,q} \text{Stock Ret}_{j,t\to t+1} \times \mathbf{1}\left(\text{Stock Ret}_{j,t\to t+1} < 0\right) \Big] + \Gamma X_{j,t} + \epsilon_{j,t} \end{aligned} \tag{IA.41}$$

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.[33] In appendix Figure IA.6 I re-estimate (7) 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 this type of compensation is not a major driving force behind my findings; the consistent

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

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

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.

Finally, in appendix Table IA.12 I replace raw excess stock returns with the same alternative labor demand shifters as in appendix Table IA.8, but now I allow estimates to vary for high- and low-productivity firms. Both for brevity and because one of my alternative labor demand shifters (customers' excess stock returns) is only available for a subset of firms, here I only sort firms into above- and below-median labor productivity. All these measures (besides perhaps the Fama-French stock return residual) are likely to be less naturally correlated with the value of equity-based pay than in my baseline specification, and I continue to find consistent significantly lower elasticity estimates—with very similar estimates quantitatively—among productive firms using these alternative shocks to labor demand. The use of these other demand shifters also helps alleviate concerns that differential exposure to confounding unobserved labor supply shocks cause productive firms to appear to face less elastic supply curves.

IA.5 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.7 demonstrates that any of a number of different reasonable proxies to difference out common market shocks yields extremely similar estimates of supply elasticities, suggesting that the way I control for these shocks is not of first-order quantitative importance (see section IA.3 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,jth entry as follows:

$$A_{i,j,t} = \log\left(1 + \sum_{\tau=t-2}^{t} (R_{i,j,\tau} + L_{i,j,\tau})\right) \tag{IA.42}$$

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

$$1 - \frac{\sum_i A_{i,j,t} \times A_{i,j',t}}{\sqrt{\sum_i A_{i,j,t}^2}\sqrt{\sum_i A_{i,j',t}^2}}, \tag{IA.43}$$

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 but still allow entries to be increasing in the number of workers that flow between the two 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 $k = 10$ 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.6 Simple Example: Comparing Firm Value in Monopsony and Counterfactual Competitive Equilibria

To see how firm value might change in a competitive equilibrium relative to the monopsony equilibrium, consider the following simple setting. There is a representative firm which solves the following:

$$V = \max_w AL(w)^{1-\alpha} - wL(w) \tag{IA.44}$$

where $L(w) = w^{\varepsilon}$ and ε is the supply elasticity. The firm's optimal wage and employment are

$$w = \left(\frac{\varepsilon}{\varepsilon+1}(1-\alpha)A\right)^{1/(1+\alpha\varepsilon)}, \qquad L = \left(\frac{\varepsilon}{\varepsilon+1}(1-\alpha)A\right)^{\varepsilon/(1+\alpha\varepsilon)} \tag{IA.45}$$

And the value of wage markdowns as a share of operating income is

$$\text{markdown share} = \frac{\left(\frac{\epsilon+1}{\epsilon}-1\right)w^{\varepsilon+1}}{V} \tag{IA.46}$$

Consider the calibration $A = 1$, $\varepsilon = 2.5$, and $\alpha = 0.3$. This yields an aggregate labor share of about 0.5 and a markdown share of operating income of 40%, both of which are very close to their empirical counterparts. The firm value in this calibration is $V = 0.25$.

Now suppose that the firm takes wages as given and equilibrium wages are determined by the same labor supply function:

$$V^c = \max_L AL^{1-\alpha} - wL \tag{IA.47}$$

where $L = w^{\varepsilon}$ gives the market clearing condition for wages. The resulting competitive wages and employment are

$$w^c = ((1-\alpha)A)^{1/(1+\alpha\varepsilon)}, \qquad L^c = ((1-\alpha)A)^{\varepsilon/(1+\alpha\varepsilon)} \tag{IA.48}$$

Both the wage and employment increase in the competitive equilibrium. In the same calibration as before with $A = 1$, $\theta = 2.5$, and $\alpha = 0.3$, the competitive firm value is $V^c = 0.21$. So there is a 16% reduction in firm value in the counterfactual competitive equilibrium, even though wage markdowns are worth 40% of firm income in the equilibrium where the firm takes advantage of its labor market power.

IA.7 Discussion of Alternative Explanations

IA.7.1 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.8 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

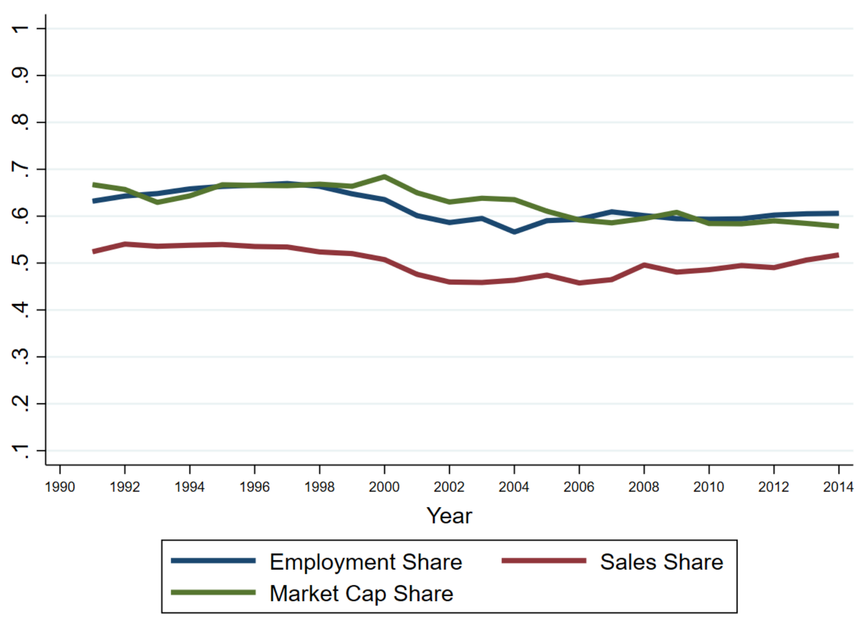

Panel B: Matched Sample Shares for Non-Financial and Non-Utilities Industries

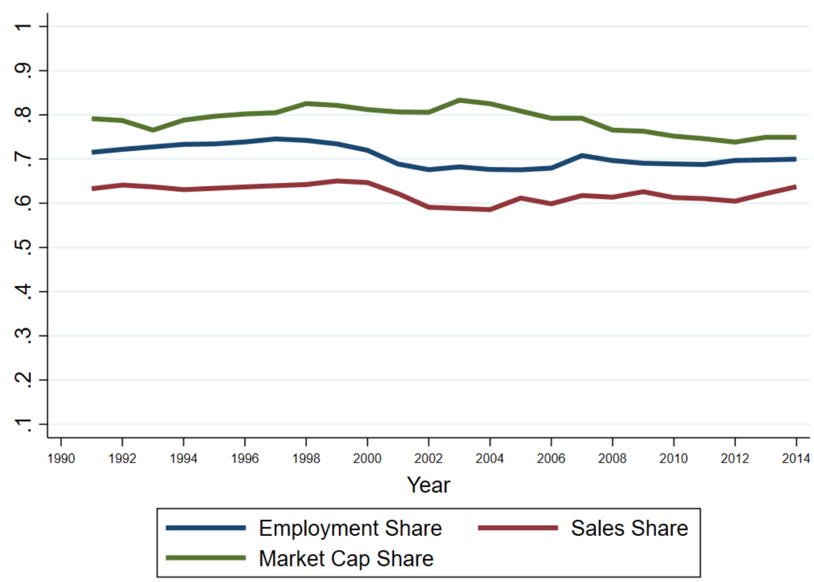

Note: This figure gives the shares of total Compustat sales, employment, and stock market cap that are represented in my Compustat-LEHD matched sample. The sample period spans 1991-2014. In panel A compares the matched sample shares to the entire Compustat universe; panel B compares the matched sample shares relative to all Compustat firms in non-financial (2-digit NAICS code 52) and non-utilities (2-digit NAICS code 22) industries, since these industries are dropped from the sample.

Figure IA.2: Aggregate Labor Shares Over Time

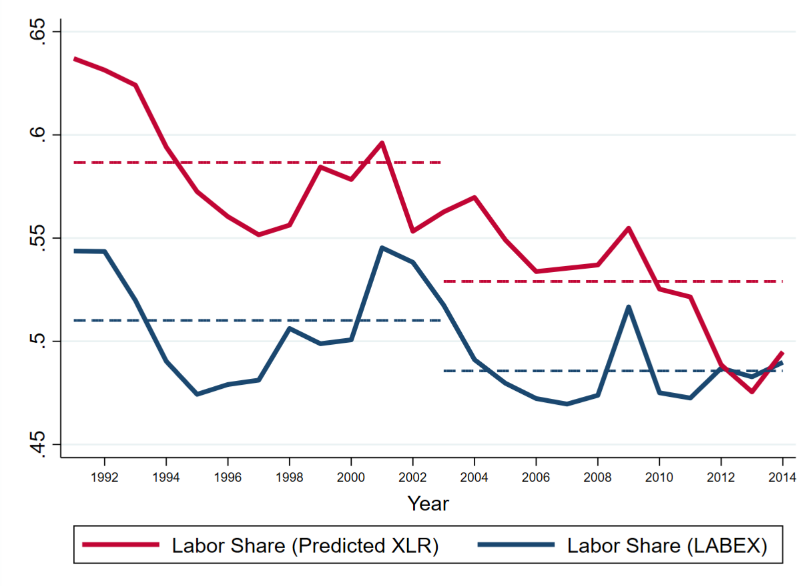

Note: This figure shows the aggregate labor shares for firms in my Compustat-LEHD matched sample. The dashed lines give the mean aggregate labor share for the given measure over the 1991-2002 and 2003-2014 subperiods. I report aggregate labor shares using two measures of firm labor expenses. Labor Share (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 valued-added following Donangelo et al. (2019). See section 4.2 in main text for further details.

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

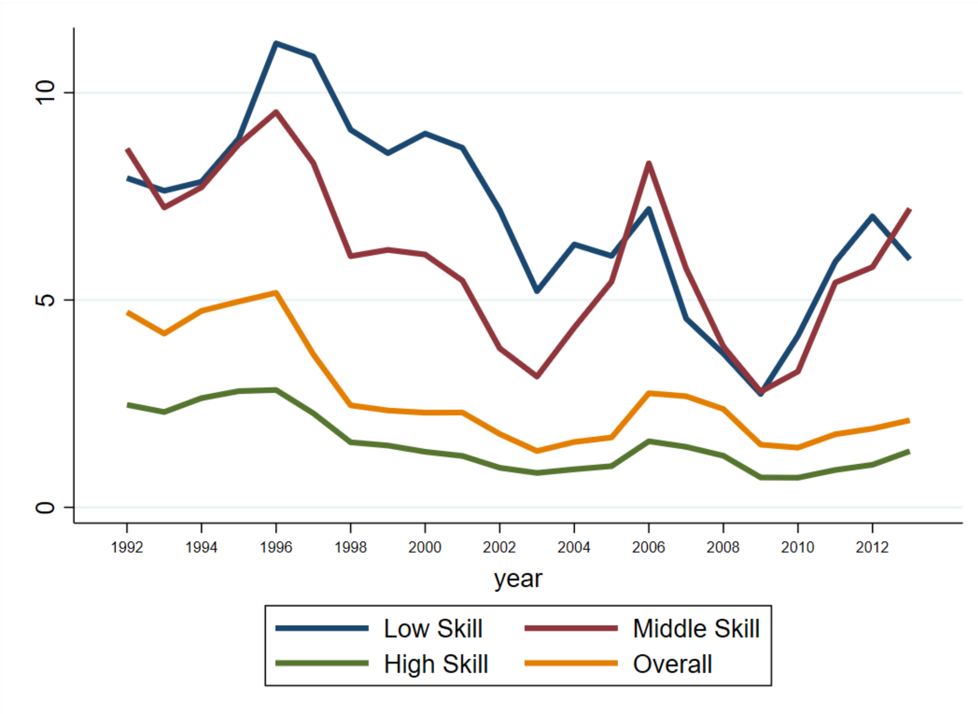

Note: This figure shows estimates of supply elasticities implied by estimating employment and wage responses from (5) for workers of different skill levels and for overlapping moving 3-year windows. The elasticity for a given year is for the moving window centered at that year (except for the start and end years, which are respectively the first or last year of 2-year window). 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 IA.4: 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.5: Spread in Average Worker Skill Between Productive and Unproductive Firms Trends Upward

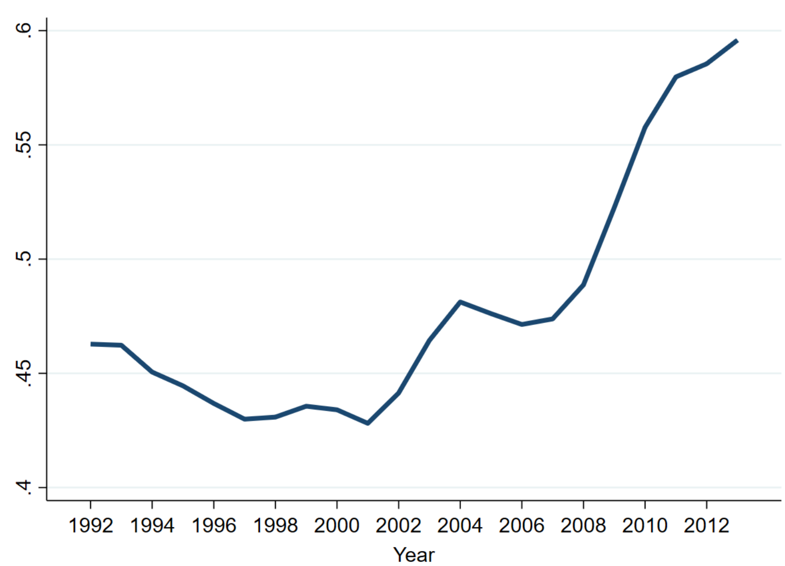

Note: This figure shows the differences for the average firm worker skill level (from (A.8) in the text) between firms in the top- and bottom-quartiles of labor productivity. The sample period spans 1991-2014.

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)

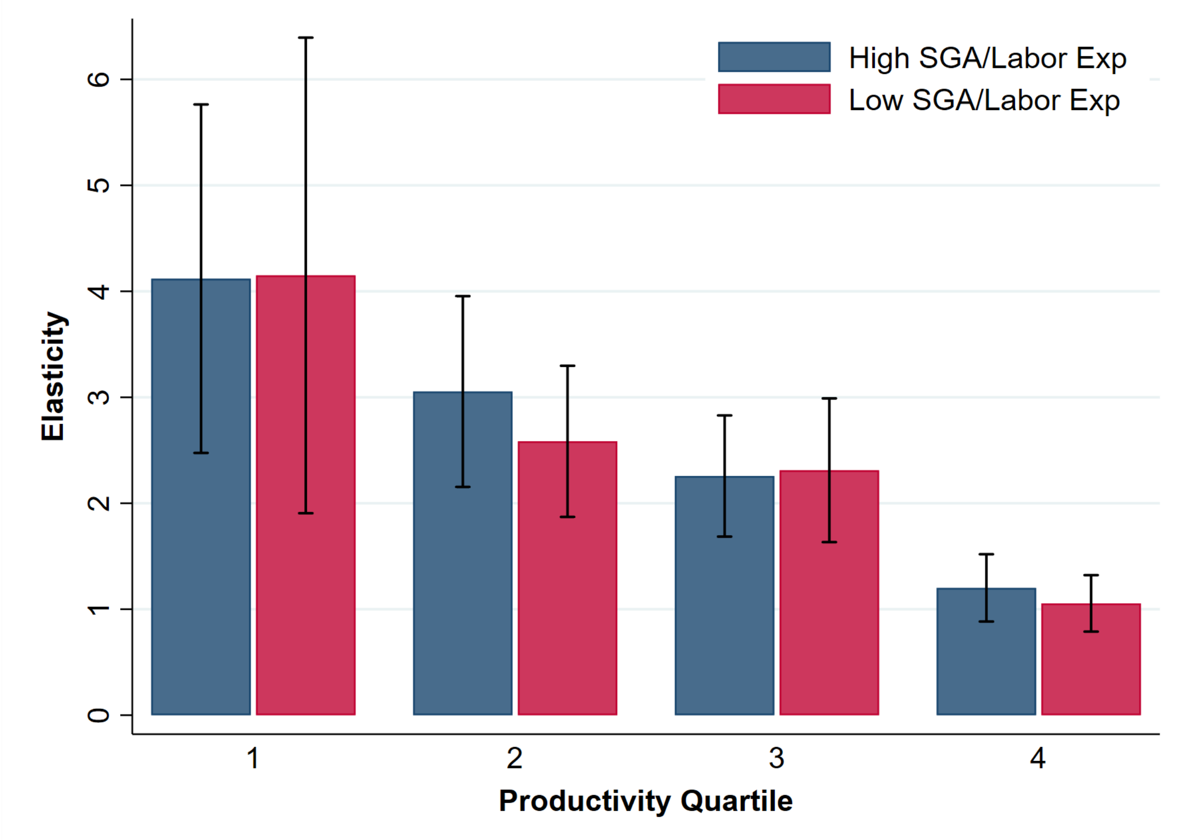

Note: This figure shows supply elasticity estimates from estimating a version of (7), where firms are double-sorted on the ratio of sales, general, and administrative expenses to LEHD labor expenses and labor productivity. Following Eisfeldt et al. (2021), this measure is used as a proxy for intensity 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 (Eisfeldt and Papanikoloau, 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

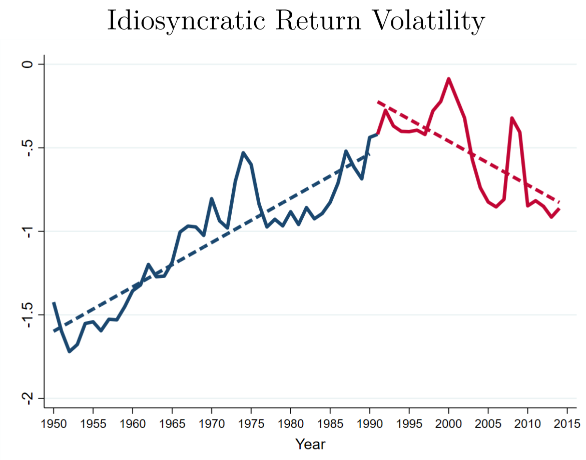

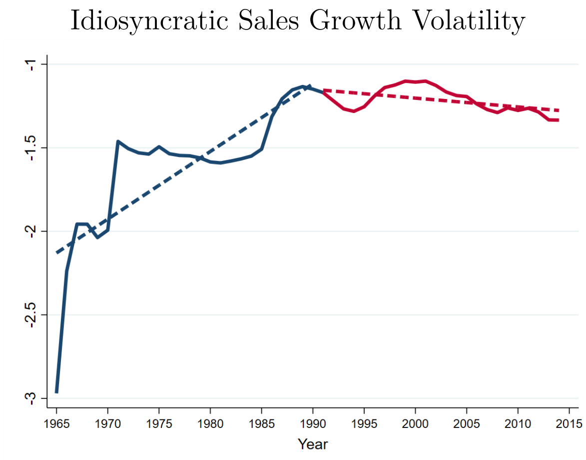

Note: This figure plots the log average idiosyncratic stock return and sales growth volatility from Hartman-Glaser et al. (2019). The data span 1950-2014 for returns and 1965-2014 for sales. The 1991-2014 period covered in my paper is shown in red and the period before is shown in blue. Dashed lines 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: Productive Firms Hire More Skilled Workers on Average

Dep Var: Log Firm Average AKM Worker Effects				
log VA/Worker	0.189	0.147	0.163	0.110
	(0.011)	(0.018)	(0.011)	(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.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
			Whole Firm				
Employment	0.12	0.09	0.12	0.11			
	(0.01)	(0.01)	(0.01)	(0.01)	0.249	0.13	43500
Wages	0.03	0.03	0.05	0.09			
	(0.003)	(0.003)	(0.004)	(0.01)	0.000	0.42	43500
Elasticity	4.32	2.71	2.34	1.17			
	(0.87)	(0.43)	(0.28)	(0.15)	0.001		
			Low Skill Workers				
Employment	0.13	0.10	0.14	0.14			
	(0.01)	(0.01)	(0.01)	(0.02)	0.522	0.11	43500
Wages	0.01	0.02	0.02	0.03			
	(0.002)	(0.003)	(0.003)	(0.003)	0.000	0.29	43500
Elasticity	10.51	6.37	6.77	4.66			
	(3.04)	(0.90)	(1.08)	(0.39)	0.042		
			Middle Skill Workers				
Employment	0.12	0.08	0.12	0.10			
	(0.02)	(0.01)	(0.01)	(0.01)	0.307	0.09	43500
Wages	0.01	0.02	0.02	0.03			
	(0.002)	(0.002)	(0.002)	(0.003)	0.000	0.18	43500
Elasticity	8.20	5.30	5.30	3.55			
	(1.70)	(0.79)	(0.71)	(0.44)	0.014		
			High Skill Workers				
Employment	0.10	0.08	0.11	0.09			
	(0.01)	(0.01)	(0.01)	(0.01)	0.148	0.08	43500
Wages	0.05	0.06	0.08	0.12			
	(0.01)	(0.004)	(0.01)	(0.01)	0.000	0.45	43500
Elasticity	2.04	1.39	1.39	0.74			
	(0.38)	(0.23)	(0.17)	(0.11)	0.004		

Note: This table contains supply elasticity estimates for firms sorted on log value-added/worker quartiles as in (7) 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 2 in main text for more details.

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

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

$$y_{j,t} = \alpha_{q(j,t)} + \alpha_{I(j),t} + \epsilon_{j,t} \tag{IA.49}$$

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.

Table IA.6: Model Versus Data Moments by Subperiod

Moment	1991-2002 Model	1991-2002 Data	2003-2014 Model	2003-2014 Data
Separations Rate	0.36	0.32	0.29	0.27
Incumbent Premium	0.13	0.13	0.17	0.17
Inc/Rec Wage Pass. Ratio	1.67	2.10	2.11	1.59
Elasticity (Overall)	2.48	2.92	1.66	1.88
Elasticity (Inc.)	1.82	2.04	1.21	1.20
Elasticity (Rec.)	5.68	8.17	5.13	5.99
Elasticity (Q1)	5.04	4.86	3.16	3.25
Elasticity (Q2)	2.84	3.00	1.93	2.18
Elasticity (Q3)	2.20	2.95	1.51	1.55
Elasticity (Q4)	1.65	1.34	1.14	0.85
Log Labor Share	-0.49	-0.49	-0.52	-0.53
Log Labor Share (Q1)	-0.35	-0.21	-0.31	-0.23
Log Labor Share (Q2)	-0.46	-0.40	-0.48	-0.38
Log Labor Share (Q3)	-0.53	-0.50	-0.57	-0.53
Log Labor Share (Q4)	-0.64	-0.90	-0.72	-0.98

Note: This table compares model and data moments for my subsample calibrations of the dynamic wage-posting monopsony model from section 3 in the main text. "Inc" and "Rec" denote incumbents and recruits, respectively. The calibrations are made to fit data moments for the 1991-2002 and 2003-2014 subperiods instead of the full sample moments from 1991-2014. The label Q1 denotes the bottom quartile of labor productivity, Q4 denotes the top quartile, etc. The calibration of 5 model parameters explicitly targets the 7 moments that are not labor productivity quartile-specific, while the remaining 8 productivity quartile moments are not targeted.

Table IA.5: Calibration by Subperiod

Parameter	Explanation	1991-2002	2003-2014
α	(One minus) Labor returns to scale	0.22	0.22
β	Supply Elasticity Shifter	0.5	0.34
b	Reservation Wage	0.549	0.515
σ_z	Labor productivity volatility	0.141	0.149
$\bar{C}$	Convex adjustment cost level	0.38	0.745

Note: This table gives two alternative calibrations of the dynamic wage-posting monopsony model from section 3 in the main text. The calibrations are made to fit data moments for the 1991-2002 and 2003-2014 subperiods instead of the full sample moments from 1991-2014.

Table IA.7: Labor Supply Elasticity Estimates Are Insensitive to Additional Controls for Common Market Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Elasticity	2.525	2.397	2.492	2.536	2.566	2.568	2.585	2.59	2.614
	(0.305)	(0.291)	(0.300)	(0.289)	(0.289)	(0.307)	(0.303)	(0.312)	(0.317)
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 × Year FE	X					X	X	X	X
Labor Market × Year FE				X	X*	X	X*	X	X

Note: This table shows estimates of (5) in the main text when different sets of controls are considered. 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.5 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.3.1 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.8: Baseline Supply Elasticity Estimates are Robust to Alternative Shocks to Labor Demand

	Customer Ret	FF Resid	Patents	Earnings Ret
Elasticity	2.01	2.34	2.77	2.67
	(0.53)	(0.29)	(1.11)	(0.42)
Baseline Elasticity	2.12	2.53	2.53	2.53
	(0.33)	(0.30)	(0.30)	(0.30)
P-Value (Diff)	0.79	0.25	0.83	0.55
N	11500	45000	45000	44500

Note: This table gives supply elasticity estimates using different labor demand shifters. All specifications have the baseline controls from (5) in main text, including industry $\times$ year fixed effects. "Customer Ret" uses the stock returns of the firm's customers instead of the firm itself. "FF Resid" uses cumulative log abnormal returns relative to the Fama-French 5-factor model augmented with momentum. "Patents" uses patent induced shocks from Kogan et al. (2020). Finally, "Earnings Ret" follows Daniel et al. (2019) in using stock returns in excess of the market return in a 4 day window around earnings announcement instead of stock returns. Elasticities computed as the ratio of the employment and wage responses to the given shock. The baseline elasticity is the elasticity estimate for the my main method using annual excess returns as a labor demand shock. In the first two columns I report baseline elasticities for the selected subsample for which the given shock is available. "P-Value (Diff)" gives the p-value on the differences between the elasticity using the given shock and for my baseline estimate. See appendix section IA.3 for further details. 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.9: 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	0.10	0.11	0.11			
	(0.01)	(0.01)	(0.01)	(0.01)	0.641	0.13	43500
Wages	0.03	0.04	0.05	0.09			
	(0.004)	(0.003)	(0.004)	(0.01)	0.000	0.42	43500
Elasticity	4.03	2.62	2.24	1.23			
	(0.79)	(0.38)	(0.28)	(0.19)	0.001		
			Low Skill Workers				
Employment	0.12	0.12	0.13	0.14			
	(0.01)	(0.01)	(0.02)	(0.01)	0.285	0.11	43500
Wages	0.01	0.02	0.02	0.03			
	(0.003)	(0.003)	(0.002)	(0.004)	0.016	0.29	43500
Elasticity	8.79	6.84	6.66	4.95			
	(2.12)	(1.69)	(1.06)	(0.64)	0.122		
			Middle Skill Workers				
Employment	0.11	0.10	0.11	0.10			
	(0.01)	(0.01)	(0.01)	(0.01)	0.444	0.08	43500
Wages	0.02	0.02	0.02	0.03			
	(0.002)	(0.002)	(0.002)	(0.003)	0.005	0.18	43500
Elasticity	7.28	5.28	5.74	3.58			
	(1.43)	(0.87)	(0.88)	(0.46)	0.031		
			High Skill Workers				
Employment	0.10	0.09	0.09	0.09			
	(0.01)	(0.01)	(0.01)	(0.01)	0.266	0.08	43500
Wages	0.05	0.06	0.08	0.12			
	(0.01)	(0.004)	(0.01)	(0.01)	0.000	0.45	43500
Elasticity	2.05	1.46	1.25	0.74			
	(0.35)	(0.20)	(0.14)	(0.13)	0.002		

Note: This table contains supply elasticity estimates for firms sorted on log value-added/worker quartiles as in Table IA.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 2 in main text for more details.

Table IA.10: 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	0.17	0.20	0.19			
	(0.02)	(0.02)	(0.02)	(0.02)	0.563	0.20	34000
Wages	0.03	0.04	0.05	0.08			
	(0.002)	(0.002)	(0.004)	(0.01)	0.000	0.51	34000
Elasticity	7.78	4.79	4.42	2.44			
	(0.88)	(0.54)	(0.63)	(0.34)	0.000		
			Panel B: Wage Growth Using Changes in AKM Firm Effects				
Employment	0.12	0.09	0.12	0.11			
	(0.01)	(0.01)	(0.01)	(0.01)	0.240	0.13	43500
Wages	0.02	0.02	0.03	0.04			
	(0.002)	(0.002)	(0.002)	(0.005)	0.000	0.25	43500
Elasticity	6.95	4.24	4.07	2.47			
	(1.15)	(0.48)	(0.45)	(0.37)	0.003		
			Panel C: TFP Sort				
Employment	0.11	0.09	0.10	0.12			
	(0.01)	(0.01)	(0.01)	(0.01)	0.482	0.14	37500
Wages	0.03	0.04	0.05	0.08			
	(0.004)	(0.004)	(0.01)	(0.01)	0.000	0.43	37500
Elasticity	3.21	2.19	1.91	1.42			
	(0.64)	(0.33)	(0.26)	(0.21)	0.007		
			Panel D: LBD Employment Growth				
Employment	0.09	0.08	0.10	0.09			
	(0.01)	(0.01)	(0.01)	(0.01)	0.683	0.08	40000
Wages	0.03	0.03	0.05	0.09			
	(0.003)	(0.003)	(0.004)	(0.01)	0.000	0.42	43500
Elasticity	3.10	2.55	1.85	1.01			
	(0.75)	(0.33)	(0.28)	(0.17)	0.01		

Note: This table shows robustness checks for the elasticity estimates obtained from estimating variants of (7) in the main text. All equations use the baseline set of controls from (7), which includes industry × year fixed effects. Panel A estimates a variant of (7) 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 İmrohoroğlu 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 2 in main text for more details.

Table IA.11: 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	0.037	0.056	0.092
	(0.006)	(0.006)	(0.007)	(0.010)
Wages (Ex Ret < 0)	0.021	0.024	0.040	0.089
	(0.005)	(0.007)	(0.005)	(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.41). See appendix section IA.4.2 for further details.

Table IA.12: Productive Firms Have Lower Supply Elasticity Estimates when Using Alternative Shocks to Labor Demand

Productivity:	Below Median	Above Median	P-val (Above - Below)
Elasticity Estimates—Customer Sample			
Customer Ret	3.40	1.45	
	(1.47)	(0.52)	0.106
Baseline	2.86	1.47	
	(0.55)	(0.24)	0.008
Elasticity Estimates—Main Sample			
FF Resid	3.35	1.41	
	(0.53)	(0.18)	0.000
Patents	2.99	1.97	
	(0.65)	(0.69)	0.038
Earnings Ret	3.34	1.87	
	(0.66)	(0.29)	0.011
Baseline	3.53	1.67	
	(0.57)	(0.17)	0.000

Note: This table shows supply elasticity estimates using alternative labor demand shifters as in appendix Table IA.8, except now firms are sorted into above- and below-median labor productivity bins. See appendix section IA.3.2 for further details on these variables. All specifications have the baseline controls from (5) in main text, including industry × 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 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: 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	1.64	1.72
	(0.33)	(0.32)	(0.34)
P-Val (Diff)		0.57	0.35
N	4200	4200	4200
Baseline Controls	X	X	X
Union Dummies		X	X
Union Dummies × Excess Return			X
Industry × Year FE	X	X	X
Union Win Excess Return	-0.08***		

Panel B: Non-Compete Changes

	(1)	(2)	(3)
Elasticity	2.37	2.36	2.4
	(0.50)	(0.50)	(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 × Excess Return			X
Industry × Year FE	X	X	X

Note: Panel A of this table shows how elasticity estimates change when including dummies for a union election occuring 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.