

Firms with Benefits? Nonwage Compensation and Implications for Firms and Labor Markets*

Paige Ouimet[†]

Geoffrey A. Tate[‡]

Current draft: July 2023; First draft: May 2022

Abstract

Nondiscrimination regulations, administrative costs, and fairness considerations can constrain the within-firm distribution of benefits. Using novel job-level data on nonwage benefits, we find a larger firm component of the variation in benefits than wages. We show that the presence of high-wage colleagues (or high-wage peers in other local firms) positively predicts workers' benefits, controlling for job characteristics including own wages. Firms with more generous benefits attract and retain more high-wage workers, but also reduce their reliance on low-wage workers more than low-benefit peers. Our results suggest that benefits disproportionately matter for worker-firm matching and, hence, compensation inequality.

*We would like to thank Rik Sen, Lea Stern, Isaac Hacamo, Carlos Avenacio-Leon, Vicente Cunat and seminar participants at the Federal Reserve, Tulane University, the University of Miami, the University of Utah, Chinese University of Hong Kong (Shenzhen), the 2022 Labor and Finance Group Conference, the 2022 FSRDC Annual Research Conference, the 2023 RCFS Winter Conference, the 2023 Society of Labor Economists Annual Conference, the 2023 Conference on Market-Based Solutions for Reducing Wealth Inequality, the 2023 Financial Intermediation Research Society Annual Meeting, and the 2023 CSEF-RCFS Conference on Finance, Labor and Inequality for helpful comments as well as Theodore To and Maury Gittleman for help with the data. Any views expressed are those of the authors and not those of the U.S. Census Bureau or Bureau of Labor Statistics. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers, as well as officers of the Bureau of Labor Statistics, have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1589. (CBDRB-FY22-P1589-R9575; CBDRB-FY23-0155). This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

[†]Kenan-Flagler Business School, UNC. Email: Paige_Ouimet@unc.edu.

[‡]University of Maryland and NBER. Email: gtate@umd.edu.

1 Introduction

Does the firm for which someone works matter for her compensation? Classical models say no: wages are equal to a worker’s marginal product, implying that there should be no heterogeneity in wages across firms after controlling for worker skill. Contrary to this prediction, Abowd, Kramarz, and Margolis (1999) finds that up to 20% of the variance in wages can be attributed to firm fixed effects. Yet, recent work argues that Abowd, Kramarz, and Margolis (1999) attributes too large a role to the firm by failing to account for factors such as limited mobility bias (Bonhomme et al., 2023). We argue that an important missing component in this debate is the role of nonwage benefits. Non-cash compensation, such as health insurance, retirement benefits, and paid leave, has become increasingly important over time in U.S. labor markets and now comprises 30% of total compensation (Bureau of Labor Statistics, 2021). Moreover, firms are likely to limit the differences between the benefits they pay to different workers at the firm because of binding regulatory and administrative constraints. Coupled with differences in the mix of worker skills across firms, this equalization of benefits within the firm can generate differences across firms in compensation to otherwise identical workers. These cross-firm differences, in turn, have implications for the distribution of human capital. Workers are likely to have a preference for more generous firms. To the extent that equilibrium firm benefit levels are more responsive to the outside options of high-skill workers, within-firm benefits compression can also increase the cost to firms of employing low-skilled labor.

To test these predictions, we perform the first match of detailed confidential job-level benefits data from the Bureau of Labor Statistics’ (BLS) National Compensation Survey (NCS) with confidential firm- and worker-level administrative data from the U.S. Census Bureau’s Longitudinal Business Database (LBD) and Longitudinal Employer-Household Dynamics (LEHD) program. We find evidence of a larger firm component in benefits than in wages. For example, between- (within-) firm variation accounts for 86.3% (13.7%) of the variance in health benefits. By contrast, within-firm variation can explain 41% of the total variance in wages in the same sample. Using a panel regression framework that controls for occupation, location, industry,

and wage, we confirm that the presence of highly-compensated employees in unrelated divisions of a firm exerts upward pressure on health benefits. Moreover, the average health benefits paid by other firms in the firm’s headquarters location and industry predicts the health benefits paid to workers elsewhere in the firm. The resulting differences in firm-level benefits affect turnover dynamics. Higher instrumented benefits predict lower turnover, particularly among a firm’s low-wage workers. Separating entry from exit, we find that higher instrumented benefits predict more worker entry than exit at the top of the wage distribution. However, they predict the opposite among low-wage workers. Together these patterns suggest that firms actively substitute away from low-skilled workers to avoid paying them the benefits required to attract and retain their high-wage colleagues. Consistent with this interpretation, we find that high benefits are associated with lower market valuations, particularly among firms that rely on jobs that are more difficult to automate.

There are a number of economic forces that can lead to within-firm compression of benefit offerings to employees. Regulations limit the ability of firms to provide different benefits to different workers in the firm. For example, the Health Insurance Portability and Accountability Act (HIPAA) nondiscrimination regulations allow employers to charge different premiums within a firm only across distinct “bona fide employment-based classifications” such as full-time versus part-time status.¹ Within an employment-based classification, all employees must be charged similar premiums. Likewise, the Employee Retirement Income Security Act (ERISA) mandates that qualified retirement plans (i.e., plans endowed with certain tax benefits) must not discriminate by disproportionately providing benefits to high-wage employees.² Beyond regulatory constraints, firms may offer one set of benefits across the board to all employees to minimize the administrative burden of managing their plans or to promote a culture of fairness, especially among benefits such as leave that have different non-monetary value to different employees or are readily observable. Mas (2017), for example, finds evidence of wage compression following greater pay transparency, consistent with the importance of fairness considerations to

¹For more information see: <https://www.dol.gov/sites/dolgov/files/ebsa/about-ebsa/our-activities/resource-center/faqs/hipaa-compliance.pdf>.

²Similar IRS nondiscrimination rules apply for self-insured health plans.

compensation policies. Thus, overall, we predict that we should observe less within-firm variation in nonwage benefits than we observe in wages. The effect is likely to be strongest among health and retirement benefits for which firms face explicit nondiscrimination standards; however, other mechanisms can induce a degree of compression even among benefits such as leave to which such standards do not apply.

The level of benefits that a firm must provide to its workers, in turn, depends upon the markets in which the firm operates. These markets could be segmented by skill or location, and different firms operate in different combinations of markets. Thus, a firm that operates in high-benefit markets will have to pay higher benefits to attract and retain workers than a firm that does not. Together with the forces that tend to equalize benefits across workers within a firm, these cross-firm differences can generate a firm-specific component to nonwage benefit payments. Importantly, firms cannot fully undo these effects on total compensation by adjusting wages given minimum wage laws, downwards wage rigidities and differences in the salience and relative valuation of nonwage benefits. Thus, a “firm component” to worker compensation is more likely to appear in benefits than in wages. That is, if where a person works matters for her compensation, it is most likely to be because of differences in nonwage benefits.

We use novel Census-matched job-level benefits data from the NCS to test for a firm component to benefits and its effects on human capital allocation and firm value. For up to eight jobs per establishment, the NCS collects quarterly information on the cost per employee hour worked for 20 unique benefits. It also provides data on the average hourly wage for the job and its occupational classification, level, compensation form, and unionization status. We match this data to the LBD, which allows us to map establishments to the parent firm. Assuming that establishments and firms are equivalent in the NCS data is likely to misclassify some within-firm variation in benefits (or wages) as between-firm variation. We also link the firm-matched benefits data with the worker data from the LEHD program, which allows us to observe the full set of workers in each firm and their wages as well as to measure turnover dynamics.

As a first step, we run a series of variance decompositions for health insurance, retirement

and leave benefits and for wages to determine and compare the magnitudes of the within- and between-firm components. Consistent with our predictions, we find that within-firm variation explains only 13.7% of the total variance in health benefits. This finding is in sharp contrast to the pattern for wages, estimated over the same sample. There, within-firm variation explains 41% of total variance. We document a similar pattern in retirement benefits and, to a weaker extent, in leave, which is not subject to specific nondiscrimination standards. All of these patterns are robust to controlling for differences in benefits or wages across states, industries, or occupations. Moreover, we document sharply increasing between-firm variation in health and retirement benefits over time, but only a modest change in the between-firm variation in wages (and leave). On the other hand, there is a modest decline in the within variation for each type of benefit and little to no change in the within variation in wages.

Having established the presence of a firm component in benefits that exceeds the firm component in wages and grows in relative importance over time, we use a panel regression approach to test for evidence of our proposed mechanism. We focus on health benefits for this analysis because, unlike health, retirement plans often have an explicit link to worker wages that could generate mechanical correlations. Likewise, the cost of leave depends directly on the hourly wage of the employee. We use the LEHD data to identify workers from each sample firm who are employed at establishments in industries that we do not observe in the NCS data, limiting the sample to only conglomerate firms. We then test whether the presence of high-wage workers in unrelated divisions is associated with higher health benefits at other jobs in the firm. We control for job-level wages and also use industry and occupation fixed effects to isolate the comparison to workers in the same occupation and industry, but who have colleagues in other divisions of the firm with different wage profiles. Confirming our prediction, we find that the presence of high-wage workers elsewhere in the firm is associated with a roughly 11.5% increase in health benefits from its mean value.

To generate between-firm variation, this within-firm convergence in benefits must be accompanied by differences across firms in the equilibrium benefit level. To test this part of our mechanism, we first identify the headquarters location of each firm as the LBD establishment

with the highest payroll per employee. We then use the NCS data to measure the leave-one-out average health benefit among firms in the industry and MSA in which this establishment is located. Using the same regression framework, we show that a one-standard deviation increase in the leave-one-out average health benefits provided by peer firms is associated with a 4.2% increase in health benefits from its mean value. Moreover, this effect is nearly the same if we limit the sample to establishments that are not located in the same MSA as the firm's headquarters. Thus, we find strong evidence consistent with both parts of the economic mechanism that we propose to explain the firm component of nonwage benefits.

One potential alternative mechanism that could generate a firm-specific component of nonwage compensation is positive assortative matching between high-skill workers and high-productivity firms. While this mechanism and ours are not mutually exclusive, it is unlikely that sorting alone can explain our findings. Because we focus on the relative between-firm variation in benefits and wages, high-skill workers would not only need to sort into high-productivity firms, but, crucially, compensation for skill differentials would need to be concentrated in the form of nonwage benefits (and more so in health and retirement benefits than in leave benefits). Likewise, the presence of high-wage workers in unrelated divisions of a firm would need to capture skill differentials that must be compensated via benefits, even after controlling directly for wages. To distinguish between stories, we show that both the presence of high-wage workers in the firm and the benefits paid by other firms in the headquarters location have, if anything, a larger positive association with the benefits in low-wage jobs than with benefits in high-wage jobs, which is consistent with our mechanism, but contrary to the assortative matching story. We also directly control for labor productivity and show that our findings are robust.

Next, we consider the implications of firm-level differences in benefits for worker turnover. For a broad class of job ladder models, a shift in compensation leads to a decline in employee turnover and quits (e.g., Dube, Lester, and Reich (2016)). Consistent with this prediction, we find that a one standard deviation increase in the average health benefits at an establishment is associated with a 15% decrease in worker turnover from its mean value. Despite our rich set of fixed effects and firm-, establishment-, and job-level controls, endogeneity could still

cloud the interpretation of this relation. For example, unobserved firm quality could drive both average benefits and worker turnover. To address this concern, we exploit the variation in establishment benefits as a result of differences in the industry and location in which the overall firm's headquarters is located. Instrumenting for establishment benefits, we confirm that benefits reduce worker turnover. The required exclusion restriction for this instrument could be violated if the benefits paid by other firms in the headquarters location and industry directly affect turnover at the establishment. This channel would predict a positive relation with turnover rather than the negative relation that we find. Nevertheless, we restrict our sample only to establishments that are located in different MSAs from the firm's headquarters to break the potential direct link. Despite the reduced power of the tests, we find similar results.

We also test for differences in the turnover effects in different parts of the wage distribution. Because fixed nonwage benefits make up a larger portion of total compensation for low-wage workers than for high-wage colleagues, a reasonable conjecture is that high establishment-level benefits will have a stronger negative effect on turnover among low-wage workers. We confirm this conjecture in the data. We also break turnover into its entry and exit components. A firm that provides high nonwage benefits to attract and retain high-skilled workers has an incentive to shed low skill workers because it will also have to provide them the elevated firm-wide benefits. We find evidence consistent with this prediction. Higher establishment-level benefits lead to a higher likelihood that entry into the establishment exceeds exit among high-wage workers, but has the opposite effect among low-wage workers. This pattern is consistent with high benefit firms shedding low-wage employees, but not replacing them with new hires. High labor costs for low-skill workers due to benefit compression within the firm increases the attractiveness of investments in labor-saving automation or outsourcing of low-skill tasks. Thus, our results can link the need to pay increasingly high benefits throughout the firm to attract and retain high-skill workers with a trend toward specialization by firms in high- or low-skill tasks. This effect can magnify the positive correlation we observe in the data between wages and nonwage benefits. We find that measures of inequality using just cash compensation understate compensation inequality. The ratio of total compensation at the 90th percentile

(including benefits) to total compensation at the 10th percentile is 9% greater than the same ratio calculated using only wages.

As a final step, we link firm-level variation in benefits with market valuation ratios. Consistent with the costs of providing generous benefits required by high-skill workers to all workers in the firm, we find a negative association between benefits and value. The relation is most pronounced among firms that employ workers in manual jobs (low-skill manual jobs are difficult to automate), but less among firms that employ more workers in abstract tasks (abstract occupations tend to require high skill). Notably, the association between wages and value is positive, suggesting it is frictions associated with benefit payments and not inefficiencies in setting total compensation that drive the relation.

Our results contribute to a number of literatures. First, we provide a new angle on the debate about the role of the firm in determining worker compensation. Our baseline analysis builds on the literature that studies between- and within-firm decompositions of wages. Song et al. (2019) uses IRS Masterfile data to show that within-firm variance is higher relative to between-firm variance for wages, but that between-firm variance has increased more over the past thirty years. Barth et al. (2016) documents a similar pattern at the establishment level using Census data. Haltiwanger, Hyatt, and Spletzer (2022) finds that much of the rising between-firm variance in wages is explained by rising between-industry dispersion in wages. We provide a novel extension of this approach to nonwage benefits, finding greater increases in between-firm variation in benefits than we observe in wages. Kristal, Cohen, and Navot (2020) estimates a variance decomposition of nonwage benefits between and within establishments (rather than firms). Observing variation at the firm-level is important given IRS nondiscrimination tests must be evaluated across all employers with sufficient common ownership or activities. Kristal, Cohen, and Navot (2020), instead, is primarily interested in measuring inequality in nonwage benefits and thus makes a number of different methodological decisions in addition to the unit of analysis, such as including part-time jobs. We are interested in differences in compensation rates for equivalent employees across firms and, hence, limit the sample to full-time jobs. Given the different objectives of our paper, we also go beyond the analysis in Kristal, Cohen, and Navot

(2020), documenting the mechanism behind the (larger) firm-specific variation we identify in benefits (than in wages) and showing its implications for turnover, quits, and labor demand.

We also contribute to the literature on the distribution of nonwage benefits. Oyer (2008) finds a positive correlation between benefits and firm size and Freeman (1981) finds that unionized jobs provide higher nonwage benefits. Liu et al. (2022) measures correlations between employee ratings of nonwage benefit categories from Glassdoor and firm-level characteristics. A related literature seeks to estimate more generally the value employees place on nonpecuniary amenities, such as job satisfaction, typically taking a more structural approach. Sockin (2022) studies amenities ratings from Glassdoor, finding that they are positively correlated with wages. Lamadon, Mogstad, and Setzler (2022) uses a structural approach to estimate unobserved job amenities and argues that more productive firms offer greater amenities. Dube, Naidu, and Reich (2022) documents a positive correlation between wages and nonwage amenities using a survey experiment at Walmart. We contribute by providing novel correlations between firm characteristics and actual employer benefit costs for a large sample of U.S. firms.

We also build on the literature, including Mitchell (1982) and Buchmueller and Valletta (1996), that documents cross-sectional correlations between the availability of health insurance and turnover. Many of these papers are limited in their ability to control for other job characteristics that are correlated with the availability of health insurance, hampering the ability to assign causality. Madrian (1994) improves upon these approaches by using the availability of spousal health insurance to better identify a causal relationship between availability of health insurance and lower turnover. However, her analysis focuses on worker-level job choices (labor supply) rather than the implications of benefit-setting practices for the firm (labor demand). Cole and Taska (2023) uses both a reduced form IV and structural approach to study the value employees place on defined contribution retirement plans, showing that a typical firm could improve the probability of job acceptance by substituting a more generous employer match in their 401(k) for salary payments. We contribute by estimating the relation between the dollar costs of health benefits and employee attraction and retention, both in aggregate and for employees in different parts of the wage distribution. Because of the unique NCS-Census

match, we can exploit granular market level differences in average health benefits to address endogeneity of firm-level benefit costs. In contemporaneous work using a different sample and estimation approach, Gao et al. (2023) study the effect of health insurance costs on firm employment. Their findings complement our worker-level analysis. We argue that frictions which inhibit firms' ability to provide differential benefits to high- and low-wage workers create an incentive for firms to specialize employment in one part of the wage distribution, for example by outsourcing employment of necessary low-wage workers to other firms. Consistent with this mechanism, Gao et al. (2023) show that shocks to health insurance costs that disproportionately increase the cost of employing low-wage workers cause firms to decrease employment of low-wage workers and increase capital investment.

Our findings also have implications for the literature on income inequality. The existing inequality literature has primarily focused on wages, because of limited data on nonwage benefits.³ Ignoring a significant fraction of total compensation is harmless if nonwage benefits mirror wages; however, we show that this is not the case. Accounting for nonwage compensation further strengthens the argument in Song et al. (2019) that dispersion in wages between firms is driving national trends, but adds nuance to what is driving the differences between firms in compensation. We document a cost to firms of having both high- and low-wage workers, predicting the pattern of increasing segregation found in Song et al. (2019). More generally, regulations that act to reduce discrimination in benefits within a firm can have the unintended consequence of incentivizing high-wage firms to replace their low skill workers through outsourcing or greater technological adoption as in Autor, Levy, and Murnane (2003), Autor and Dorn (2013) and Goldschmidt and Schmieder (2017).

Finally, and more generally, our paper relates to the literature investigating firm-level compensation practices and their implications for value. Ellul, Pagano, and Schivardi (2018) finds

³The exceptions in the literature have primarily focused on one non-pecuniary job characteristic at a time, such as undesirable work hours in (Hamermesh (1999)) and labor rights violations as in (Marinescu, Qiu, and Sojourner (2021)) or indirectly estimate the value of nonwage job amenities (job characteristics that include the nonwage benefits discussed in the paper as well as other job characteristics such as job security) using a willingness-to-pay estimate extracted from surveys as in Maestas et al. (2023) and Dube, Naidu, and Reich (2022) or job switchers as in Taber and Vejlín (2020), Lamadon, Mogstad, and Setzler (2022) and Sockin (2022).

that family firms provide lower wages in exchange for providing greater employment protections. Berk, Stanton, and Zechner (2010) argues that firms with more leverage must pay higher wages and Graham et al. (2022) finds that wages fall when firms enter bankruptcy, especially in thin labor markets or at small firms. Silva (2010) documents evidence of wage convergence in conglomerates. Our results imply similar and, perhaps, more pronounced convergence in benefits within conglomerates. To the extent that constraints on benefits are not undone with adjustments to wages, we potentially provide an alternative cost-based channel that can link the distribution of human capital in conglomerates with a diversification discount (Tate and Yang, 2015). Similarly, our results can provide a human-capital channel to explain differences in firm performance that depend on firm geography (Garcia and Norli, 2012). A related literature has looked at the implications to firms of shocks to minimum wages. Hau, Huang, and Wang (2020) find that when minimum wages increase, firms engage in capital deepening, leading to lower employment and higher capital expenditures. Our finding that high benefit firms are more likely to shed low wage workers expands upon this literature by documenting that voluntary actions by firms, as measured by relative generosity in benefits, have implications for the distribution of workers within the firm and sheds light on the cross-sectional distribution of within-firm inequality. In particular, wage inequality tracks total compensation inequality most poorly for the lowest deciles of the distribution.

2 Economic Mechanisms and Empirical Predictions

We consider two forces that can generate a firm component to benefits: (1) a tendency for firms to limit variation in benefit offerings across workers, weakening the link between worker or job characteristics and benefit costs and (2) differences in the equilibrium benefits offered by firms because of differences in the (segmented) labor markets in which they operate. These forces predict a firm component to worker benefits that is distinct from the firm component of wages studied in prior literature (e.g., Abowd, Kramarz, and Margolis (1999)). We develop several empirical predictions for our analysis based on these underlying economic mechanisms.

Firms are likely to limit cross-worker variation in benefits for a number of reasons. Providing compensation in the form of benefits rather than wages can create tax benefits both for workers and the firm; however, qualifying for those tax benefits requires firms' benefit plans to meet certain nondiscrimination standards. For example, a 401(k) sponsor must annually pass the Actual Deferral Percentage (ADP) test and the Actual Contribution Percentage (ACP) test. The ADP and ACP tests look for discrimination by comparing average deferrals and contributions of highly compensated employees to non-highly compensated employees. IRS nondiscrimination rules for health plans are similar, but only apply to self-insured health plans and arrangements in which employees contribute pre-tax dollars to pay for certain employee benefits, typically referred to as Section 125 or cafeteria plans. Employer-sponsored health insurance plans must also comply with other federal laws that ban discrimination, such as the Americans with Disabilities Act (ADA), the Age Discrimination in Employment Act (ADEA), and the Health Insurance Portability and Accountability Act (HIPAA), all of which limit the ability to design health plans that have different eligibility requirements, contributions, or benefits across workers within a firm. In the Appendix, we provide more details on these restrictions. Though nondiscrimination provisions do not apply to all categories of benefits (notably, there are no explicit standards for leave policy), they will tend to suppress within-firm variation in benefit costs in a way that does not apply to wage payments.

In addition to explicit regulatory constraints, administrative costs can exert pressure on firms to standardize benefit offerings across workers. Health plans, for example, typically include complex and multi-dimensional bundles of provisions, each of which, unlike wages, is likely to be valued differently by different subsets of workers. Increasing the number of offered plans in response would not only increase the firm's administrative burden, but could also limit the ability of the firm to offer more attractive terms by limiting the plans' economies of scale. Variation across workers in plan terms could also create costs for the firm if employees value fairness. In the context of wages, Mas (2017) finds that greater mandated pay transparency leads to a decrease in wage dispersion among city managers in California, consistent with the importance of fairness considerations to compensation policy. Even though regulatory nondis-

crimination standards do not apply to leave benefits, this mechanism could be particularly acute for leave policies because days of absence are directly observable. Together, these various channels predict more limited within-firm variation in benefits than we observe in wages.

To the degree that these factors induce unwanted convergence of total compensation across workers, the firm could conceivably undo the effects by adjusting wages. However, there are significant constraints on firms' ability to do so, given minimum wage laws, downward wage rigidities, and differences in the salience and individual relative valuations of wages and benefits. While minimum wage laws only impact a segment of employees, it is exactly employees at the lower end of the wage distribution for whom firms would most like to offset higher benefits with lower wages, assuming that their skills are least likely to be scarce in the labor market. Moreover, limited information can constrain the ability of new hires to value nonwage benefits prior to joining the firm, thus making wages relatively more salient when comparing multiple job offers. After joining, firms are likely to be limited in their ability to adjust wages lower given downward wage rigidities. This same mechanism can also limit firms' ability to attract new employees who received higher wages from their previous employer. Finally, differences in marginal tax rates will lead lower-wage workers to place less value on non-wage benefits than cash compared to higher-wage workers (Woodbury (1983)).

Variation in the firm component of benefits, in turn, depends on the labor markets in which firms operate. We focus on two sources of variation in the resulting equilibrium firm-level benefit: (1) differences in the types of workers employed and (2) differences in the benefits paid by competitors. Not all workers inside a firm will be hired in the same labor markets. Different tasks within the firm can require different non-substitutable skills. The compensation the firm must pay to workers hired to perform these different tasks will depend on the equilibrium compensation in the separate markets for workers with the requisite skills. Some of these skills are likely to be in greater supply than others. For example, an aerospace engineer with a history of patented innovation could be in more scarce supply than a data entry clerk. Then, the aerospace engineer is likely to have more market power over her human capital and to be able to extract additional rents from a hiring firm compared to the data entry clerk. These rents

can include generous benefit packages. Thus, if we consider two firms, one of which employs only data entry clerks and the other of which employs identical data entry clerks, but also employs aerospace engineers, the incentive for firms to limit cross-worker variation in benefits would predict that the data entry clerks in the second firm receive higher benefits packages than the data entry clerks in the first (and that the second firm pays higher average benefits to its workers than the first).

Geographic labor market segmentation can produce similar differences across firms in the benefit costs of otherwise identical workers. Firms may operate in markets in which the scarcity of those workers is different (stemming from differences in labor supply or demand), generating differences in bargaining power and equilibrium benefit payments. Alternatively, firms may operate in markets with very different costs of living (including, e.g., healthcare prices). These differences across locations in the (average) benefits paid to similar workers by competitor firms should be reflected in a location effect, or, alternatively, a location-industry effect if there is also within-market segmentation by worker skills. However, in the presence of a tendency for firms to equalize benefits across workers, they can also change the firm component of benefits. That is, a worker in a firm that employs workers in an unrelated division that is located in a high-benefits market will receive higher benefits than an otherwise identical worker in a firm in the same local market whose workers in other divisions are located in low-benefits markets. In general, the benefits required to attract and retain the workers in the firm with the best outside options will determine the firm component to benefit offerings.

Based on this discussion, we make a number of empirical predictions about the existence and nature of a firm component in benefits:

Prediction 1. *The fraction of the total variation in benefits that is explained by within- (between-) firm variation will be smaller (larger) than the corresponding fraction of the variation in wages that is explained by within- (between-) firm variation.*

We also assume that wages increase with worker skill level and that high-skilled workers have the highest outside options. Moreover, we assume that the greatest concentration of high-skilled

workers in the firm is at the headquarters location. Then, we have two additional predictions:

Prediction 2. *Holding job characteristics fixed, a worker's benefits will increase with the presence of highly-paid workers in unrelated divisions of the firm.*

Prediction 3. *Holding job characteristics fixed, a worker's benefits will increase with the average benefits received by workers in other firms in the same industry and local market as the firm's headquarters.*

Given a firm component to benefits that is determined in large part by the outside options of the firm's high-skilled workers, there are also implications for labor supply and demand. On the supply side, firms with generous benefit offerings will be particularly attractive to low-skilled workers, who will expect to receive higher benefits than they could obtain in the absence of high-skilled colleagues and for whom benefits make up a larger fraction of total compensation. But, high-benefit firms are also likely to demand fewer low-skilled workers given these higher costs. Instead, such firms will have a relatively greater incentive to pursue automation, outsourcing, and other labor-saving technologies to substitute for low-skilled workers in production. Moreover, to the extent that firms cannot fully specialize in the human capital distribution, outsized benefits payments to low-skilled workers are likely to be associated with lower firm valuations. Thus, we make three final predictions:

Prediction 4. *Firms with more generous benefits have lower worker turnover, particularly among low wage workers.*

Prediction 5. *Among low- (high-) wage workers, firms with more generous benefits are more likely to have worker exit (entry) rates that exceeds worker entry (exit) rates.*

Prediction 6. *Firm benefit levels negatively correlate with firm value, particularly among firms that employ workers in hard-to-automate occupations.*

3 Data

To test our empirical predictions, we use the quarterly Employer Costs for Employee Compensation (ECEC) data from the BLS’s National Compensation Survey (NCS) from the second quarter of 2004 to the fourth quarter of 2014. The data samples all civilian workers, with few exceptions.⁴ We match each establishment in the NCS data to the Census Longitudinal Employer-Household Dynamics (LEHD) data using the state unemployment insurance identifier and to the Census Longitudinal Business Database (LBD) using the Standard Statistical Establishment Listing (SSEL). For the purposes of our variance decompositions, we drop all observations with imputed data for wages or for the given benefit we are considering.⁵ As a result, our sample size varies depending on the benefit being considered. As benefits to part-time employees can differ from those provided to full-time employees, we drop all part-time jobs from our data. Finally, we drop observations in the public sector (NAICS92) and observations that cannot be matched to the Census LBD. The latter restriction is necessary because the LBD allows us to identify the firms that ultimately own each establishment included in the sample. The unit of observation in the data is an establishment-job-quarter.⁶ We have 519,000 unique observations and 11,500 unique firms in our health benefits sample, 545,000 unique observations and 12,000 unique firms in our retirement sample, and 523,000 unique observations and 11,500 unique firms in our leave sample.⁷

The BLS samples establishments using a probability proportional-to-size technique in order to produce a nationally representative sample. As a result, our sample over-represents large establishments. This is beneficial for our purposes, as our estimates are based on establishments that are economically meaningful in terms of employment. However, our results may not be generalizable to the smallest of establishments. Moreover, in order to implement this approach,

⁴Exceptions include civilian workers employed in the federal government, military, agriculture as well as unpaid and overseas workers.

⁵Imputed wage observations are relatively rare in the data. They can happen, for example, if the survey response includes non-imputed wages in quarters 1 and 3 of a year, but does not have available wage data for quarter 2.

⁶Information on sample construction in the NCS is available here: <https://www.bls.gov/opub/hom/ncs/design.htm>.

⁷Sample sizes are rounded to conform to Census conventions at the time of our disclosure request.

the BLS must first develop a frame with establishment sizes, then sample within this frame. This sampling approach is likely to undersample new establishments. Once they enter the sample, establishments are (typically) surveyed each quarter for five years before being dropped from the sample.

For each establishment in the sample, and for each quarter, the BLS estimates characteristics for up to eight jobs, with fewer jobs selected at firms with fewer employees.⁸ It is important to note that “jobs” are defined very narrowly in the data.⁹ Each job captures a worker or set of workers within a given establishment, uniquely identified within a given 6-digit standard occupational code (SOC), requiring the same level of skill and responsibility, having the same full-time or part-time status, having the same union or nonunion status, and having wages that are determined in a similar fashion (e.g. hourly or incentive-based).¹⁰ Within an establishment, jobs are sampled using a proportional-to-size approach so that jobs representing a larger fraction of total establishment employment are more likely to be sampled.

For each establishment job quarter sampled, the BLS estimates the average cost of wages and benefits for the average employee per hour worked in that job. Wages include all regular payments from the employer to employee including incentive-based pay such as commissions, production bonuses, and piece rates. We adjust all costs to 2014 dollars using the CPI. Estimations are made for five broad benefit categories: paid leave; supplemental pay; insurance; retirement; and legally required benefits, such as Social Security, Medicare, unemployment insurance and workers’ compensation. Estimates for hours-based benefits, such as paid leave, are based on wages and salaries.¹¹ Estimated costs account for non-participation by employees who would be eligible to participate. We focus on health insurance, retirement plans (defined

⁸The BLS samples up to 4 jobs at establishments with less than 50 employees, 6 jobs at establishments with employment between 50-249 and eight jobs at establishments with 250 or more employees.

⁹BLS uses the terminology “hits” instead of job. However, for ease of exposition we keep to the more common term of job.

¹⁰Workers are assigned to occupations based on the workers’ actual job duties and responsibilities, not on job titles. For a worker whose duties span more than one occupation, the occupation which involves the highest skill is assigned. If duties involve occupations of equivalent skill, then assignment is based on where more of the worker’s time is spent.

¹¹For example, for leave, the BLS estimates the number of hours of the benefit used by workers in the sampled job, multiplies this by the compensation rate, and then normalizes it by the total hours worked to get a cost per hour.

benefit and defined contribution plans) and leave (sick leave, vacation, holidays) because they are the most economically important voluntary benefits.

In Table 1, we report summary statistics for our sample of job establishment quarters. Most firms offer some form of health, retirement, and leave benefits. 87.4% of the jobs in our sample provide health benefits, 72.8% provide retirement, and 95.5% provide paid leave. Firms in our sample are more likely to provide benefits, compared to the estimates for the US economy as a whole.¹² The difference is not surprising given larger firms are more likely to offer benefits and larger firms are over-sampled in our data. There is more variation in the cost of providing these benefits. Health insurance is the most costly benefit. On average, employers in our sample pay \$3.17 per hour to provide health benefits. This compares to a national private-sector average cost of providing insurance benefits for full-time workers in 2014 of \$3.13, according to the BLS.¹³ Leave is the second most costly, with employers paying \$2.59 per hour, on average, in our data. Nationally, firms pay \$2.68 per hour on leave, as estimated by the BLS in 2014. Finally, employers in our sample pay \$1.67 on average to provide retirement benefits for their employees. This is similar to aggregate statistics from the BLS, which estimates that employers spend, on average, \$1.55 on retirement benefits in 2014.

Across our samples, the average wage ranges from \$25.22 to \$26.03, similar to the national US average of \$24.35 in 2014. On average, we find total compensation – measured as wages plus health, retirement and leave benefits – to be between \$32.61 and \$33.90. Across our time series, benefits account for 23% of total compensation. From 11.5% to 13.1% of jobs in our samples are unionized. The average job is held by between 25 and 28 employees per establishment.

We also report statistics at the establishment and firm level. As expected, establishments in our sample are large, with average employment between 981 and 1,149. Given the size of the establishments, it is not surprising that our samples include more than four jobs per establishment, on average. We observe differences in firm size across our benefit samples. The sample of firms with non-imputed leave benefits has the smallest firms, with 20,190 employees

¹²For more information, see <https://www.bls.gov/news.release/pdf/ebs2.pdf>

¹³Data is available from <https://www.bls.gov/ncs/data.htm>

on average. Firm employment in our samples with non-imputed health and retirement data is larger, at 37,050 and 44,730, respectively. These averages are similar, though larger than, the average non-financial firm in Compustat, which has 10,195 employees in 2014. The average firm in our sample has an average age of over 27 years.

In Table 2, we report correlations between benefits and establishment, firm, and job characteristics. Mirroring the approach in Table 1, we report these statistics using the sample of jobs for which health insurance benefits are non-imputed (Column 1), retirement benefits are non-imputed (Column 2), and leave benefits are non-imputed (Column 3). Given the frequency of zeros in our data (when no benefit in a category is provided) as well as the skewness in the distributions of the benefits, we transform each of the benefit costs using the inverse hyperbolic sine function. We transform employment, age and wages using the natural logarithm because zeros are never present in these variables.

We find that health, retirement, and leave benefits are positively correlated with wages. Workers who receive high wage compensation also receive high nonwage compensation. The magnitude of these correlations varies by benefit. Leave benefits and wages are highly correlated at nearly 0.75. The high correlation is not surprising given that the cost of providing paid leave is mechanically related to wages. Retirement benefits and wages have a correlation of 0.614. Both defined benefit and defined contribution benefits are often linked to wages, consistent with the high correlation. Health benefits, which are not mechanically linked to wages, have a correlation of 0.47 with wages.

Unionized jobs provide higher nonwage benefits, on average, as in Freeman (1981). Larger and older firms are associated with more generous benefits. A similar relation between wages and firm size was first documented in Brown and Medoff (1989) and between wages and firm age in Brown and Medoff (2003). Finally, we measure the correlations between benefits.¹⁴ All three benefits are positively correlated, with magnitudes between 0.52 and 0.57.

¹⁴Due to restrictions related to disclosure, when estimating the correlation between health benefits and retirement benefits using the non-imputed health sample (Column 1), health benefits are never imputed but retirement benefits could be imputed. Likewise, when estimating the correlation between retirement benefits and health benefits using the non-imputed retirement sample (Column 2), retirement benefits are never imputed but health benefits could be imputed. Mirror condition applies to the other correlations between benefits.

The positive correlation between wages and benefits suggests that income inequality is being under-estimated if only cash compensation is considered. Indeed, we confirm this prediction in Figure 1 in which we graph the distribution of both wages and total compensation, measured as wages plus health, retirement and leave benefits. Regardless of whether we use the health, retirement or leave samples, we document a steeper rise in total compensation than in cash compensation from the 10th to the 90th percentile. Interestingly, mis-measurement appears to be concentrated among the lowest income percentiles. The ratio of the 90th to 50th percentile of total compensation is nearly identical to the equivalent ratio using cash compensation. However, the ratio of the 50th to the 10th percentile of total compensation is 7.5-9% larger than the equivalent ratio using cash compensation.

4 The Firm Component of Benefits

The first step in our analysis is to test for the presence of a firm component of benefits that is distinct from the firm component of wages. We take two approaches. First, we estimate variance decompositions of benefits (and wages) into within- and between-firm components (testing Prediction 1). Second, we use a panel regression approach to test whether the variation in benefits is consistent with the underlying economic mechanisms that we propose to generate the firm effects (testing Predictions 2 and 3).

4.1 Variance Decomposition

We begin by considering decompositions of the variation in benefits. Similar to Song et al. (2019) and Barth et al. (2016), we decompose the overall variance of nonwage benefits into within- and between-firm components. However, in our context, the within-firm variation is across jobs rather than workers. Specifically, let $b_t^{j,f}$ be the nonwage benefits of job j at firm f in quarter t . The benefit can be separated into two components:

$$b_t^{j,f} = \bar{b}_t^f + [b_t^{j,f} - \bar{b}_t^f], \tag{1}$$

where \bar{b}_t^f is the average nonwage benefit for firm f in quarter t . Thus, the overall variance can be decomposed into two components:

$$var(b_t^{j,f}) = var_f(\bar{b}_t^f) + \sum_f \omega_f \times var_j(b_t^{j,f} | j \in f), \quad (2)$$

where the first term captures between-firm dispersion of firm average nonwage benefits and the second term captures the employment-weighted mean of within-firm dispersion in nonwage benefits. ω_f is the employment share of firm f in the sample. Given the frequency of zeros in our data, we use an inverse hyperbolic sine transformation of the dollar value of the given benefit per hour worked.

In Table 3, we report separate calculations for health insurance and retirement benefits (Panel A) and leave (Panel B). We compare the decomposition of nonwage benefits to a similar decomposition of wages.¹⁵ We repeat this exercise three times, using the sample of non-imputed health benefits (used for the variance decomposition of health benefits), the sample of non-imputed retirement benefits (used for the variance decomposition of retirement benefits), and the sample of non-imputed leave benefits (used for the variance decomposition of leave benefits.) To be consistent with the existing literature, we use a log transformation of wages.¹⁶

We find that 86.3% (86.0%) of the total variation in health (retirement) benefits can be attributed to differences between firms. In contrast, only 59% of the total variation in wages can be attributed to differences between firms. The smaller within-firm variation in health and retirement benefits is consistent with the regulatory and administrative costs of varying benefits across employees. We also observe that between-firm variance in health and retirement benefits is significantly larger than between-firm variance in wages. While an earlier literature has documented significant across-industry and across-firm wage differentials, we show that there is more than three times as much variation in health and retirement benefits between

¹⁵We report results using equal weighting of each observation. We also repeat the exercise using occupational weighting. The results are similar.

¹⁶If we instead use an inverse hyperbolic sine transformation of wages, the results do not vary by more than 1%. We do not report them to avoid showing duplicate results, however, they are available upon request from the authors.

firms than there is in wages.

In Panel B, we consider variation in paid leave benefits. We consider paid leave separately because the IRS and HIPAA nondiscrimination regulations that apply to health and retirement benefits do not apply to paid leave. As a result, we expect the distribution of leave benefits to be more similar to wages. However, we do not expect leave benefits to fully mirror wages. Firms may still prefer to offer more consistent leave across employees to minimize the administrative burden of managing multiple plans or out of fairness concerns. Given that coworkers can typically observe the number of days a colleague takes off, firms may face greater pressure to offer more equal leave benefits. We find higher within-firm variance in leave compared to health and retirement benefits or to wages. However, as with health and retirement benefits, we find higher between-firm variation in leave benefits than in wages. Put together, we find a distribution of the variation in leave benefits that falls between the distributions of health and retirement benefits and wages, with nearly 73% of the total variation attributed to differences between firms.

In both panels, we document lower wage variance than Song et al. (2019) and Haltiwanger, Hyatt, and Spletzer (2022). One key difference between our data and the IRS MasterFile data used in Song et al. (2019) and the LEHD data used in Haltiwanger, Hyatt, and Spletzer (2022) is that the NCS data measures hourly wage rates while the other data sources report actual annual (or annualized from quarterly) wages, which may or may not reflect an employment spell that lasts for the entire year. As a result, we are measuring different concepts of within-firm wage variance. Our measure of inequality can be thought of as a measure of intentional inequality. We capture differences in promised hourly wages for workers within and between firms, in a given quarter. In contrast, Song et al. (2019) and Haltiwanger, Hyatt, and Spletzer (2022) measure realized within and between firm variances, reflecting both differences in promised wages as well as differences in employee turnover. To the extent that lower wage workers and lower wage firms have higher rates of turnover (as we observe in Section 5), both within-firm and between-firm variance using realized annual wages will be higher than estimates using the NCS data on hourly promised wages. It is, thus, not surprising that our estimates of within

and between firm wage variance are more similar to Barth et al. (2016), which use annualized LEHD data but limits the sample to employees observed at a given firm for all four quarters.

While we argue that the main difference between our results and those in Song et al. (2019) and Haltiwanger, Hyatt, and Spletzer (2022) is driven by the difference between wage rates and realized wages, there are also other subtle differences in our data. Notably, we only consider full-time workers, which mechanically reduces within-firm variance, but also moderates between-firm variance to the extent that some firms have more part time workers than other firms. In addition, we use Census-defined firms, which account for common ownership of tax-reporting entities, as the unit of analysis while Song et al. (2019) use tax reporting entities (EINs). The two concepts typically coincide for single unit firms, but often diverge for large firms with multiple establishments.

We observe strikingly different patterns in the variance decompositions of nonwage benefits and wages. Between-firm variation accounts for a larger percentage of total variation in benefits than in wages. We attribute this difference to institutional frictions that specifically apply to the provision of nonwage benefits within a firm. However, it is also possible that differences in the distribution of nonwage benefits from the distribution of wages across geographies, industries or occupations (to the extent occupations are concentrated within certain firms) could generate similar patterns.

To address these alternative mechanisms, we adjust benefits and wages for differences across states, industries, or occupations and then recalculate the variance decompositions. Following the approach in Song et al. (2019), we first subtract the mean nonwage benefit or wage by state, occupation (6-digit SOC), or industry (6-digit NAICS). After demeaning the data, we then re-estimate the variance decomposition. We report the results in Table 4. In Column 1, we report between-firm variances in nonwage benefits after demeaning. In general, the between-firm variance is reduced, with the greatest reduction occurring with controls for industry, consistent with Haltiwanger, Hyatt, and Spletzer (2022). In Column 5, we report the same exercise for wages. We also find a general reduction, with the largest occurring with controls for occupation. Not surprisingly, demeaning by state or industry has less of an effect on within-firm variance in

benefits (Column 2) or wages (Column 6). Demeaning by occupation, which can vary within-firm, has more of an effect on within-firm variance, particularly on wages. Critically, across all three approaches to standardization and across all three benefits, we continue to observe that between-firm variation is a larger fraction of total variance in nonwage benefits than in wages. That is, after controlling for state, industry and occupational patterns in compensation, we continue to observe sharply higher differences between firms in the generosity of nonwage benefits than we do in wages.

As a final step, given the dramatic rise in wage inequality over time and especially between-firm inequality (Song et al. (2019)), we consider differences in the between-firm components of the variation in benefits in the time series. In Figure 2, we plot the variance decomposition by year starting in 2004 and extending until 2019.¹⁷ In order to include five additional years of data in our time series, we measure between and within variance at the EIN-level (a variable available in the NCS data) as opposed to the firm-level (a variable only available in the Census data and with a shorter time-series in the snapshot available at the time of our disclosure).¹⁸

In Panel A, we consider health benefits. We find a dramatic increase in the between-establishment variance in health benefits over time. In the 14 years of our sample, this component of the variance increases by 53%. The increase is more or less monotonic but appears to level out starting in 2016, which could reflect the employer mandate of the ACA that went into effect for large employers in 2015 and for smaller employers in 2016. By comparison, there is a more modest increase in the between-establishment variance in wages, increasing by only 7% over our sample period. There is a 27% decline in within-establishment variance in health benefits and no change in within-establishment wage variance.

Panel B presents similar patterns for retirement benefits. The between-establishment variance in retirement benefits increases by 50% over our sample period, while within-establishment variance decreases by 14%. On the other hand, the patterns for leave benefits are more muted,

¹⁷The variance decomposition is estimated by quarter and then averaged for each year. In 2004, this only includes quarters two to four due to missing data for the first quarter.

¹⁸We show in Internet Appendix Table 1 that the between- and within- variance decompositions are similar when using either EIN- or firm-level data.

as shown in Panel C. Between-establishment variance in leave benefits increases by 13% and within-establishment variance in leave benefits decreases by 13%. Patterns in wages are similar across all three samples. Together with the pattern in Figure 1, these results suggest that differences across firms are even more important in driving total compensation inequality than wage inequality, and that, over time, this firm component of total compensation is becoming increasingly important.

4.2 Regression Analysis

Our variance decomposition results suggest that a worker’s total compensation is determined, in part, by the firm at which she is employed. Moreover, we observe a larger firm component of the variation in benefits (health, retirement, and leave) than we do in wages. We propose institutional frictions, such as the nondiscrimination provisions in certain laws governing firm-provided benefits, as an economic mechanism generating the firm effects. Next, we test the predictions of this mechanism for the distribution of benefits across firms. First, we test whether high-wage (skill) workers within a firm exert upward pressure on the benefits provided to colleagues, controlling for their own wages (Prediction 2). Second, we test whether variation in the benefits provided by competitors in the markets in which a firm’s high-skill workers are located exerts upward pressure on benefits (Prediction 3).

4.2.1 Benefit Costs and the Within-firm Distribution of Wages

To test Prediction 2, we estimate the following linear regression specification:

$$Benefit_{jefist} = \beta HiWageExcDiv_{ft} + \mathbf{X}'_{jefist} \boldsymbol{\omega} + \delta_s + \alpha_j + \mu_i + \theta_t + \epsilon_{jefist} \quad (3)$$

$Benefit_{jefist}$ is the inverse hyperbolic sine of the dollar cost of health benefits provided to employees in job j at establishment e of firm f in industry i in state s in year-quarter t .¹⁹ We

¹⁹We use the inverse hyperbolic sine transformation instead of a log transformation because roughly 12.5% of the jobs in the sample do not provide health benefits (See Table 1).

focus our analysis on health benefits because retirement and leave benefits in our sample are often, by plan design, wage-related, while health benefits are not. Thus, there is no potential for a mechanical relation between health benefits and wages across jobs inside the firm.²⁰ $\mathbf{X}'_{\text{jefist}}$ is a vector of control variables, which may include the natural logarithms of the average hourly wage earned by workers in the job as well as firm employment, firm age, and establishment employment. δ_s is a state fixed effect, α_j is an occupation fixed effect (measured using 6-digit BLS SOC codes), μ_i is an industry fixed effect (measured using 3-digit NAICS codes), and θ_t is a year-quarter fixed effect.

To construct our independent variables of interest, $HiWageExcDiv_{ft}$, we first identify for each NCS establishment the set of LEHD establishments that are not part of the NCS sample and that operate in industries that we do not observe among the firm’s establishments that are in the NCS sample. We restrict the sample to firms for which we observe all firm divisions within the states that are part of our sample of LEHD data in order to avoid measurement error in our characterization of the wage distribution among workers in excluded divisions of the firm. We then calculate the distribution of the workers in those establishments across quintiles of the overall wage distribution (i.e., considering the full set of workers from NCS firms that we can match to the LEHD data). $HiWageExcDiv_{ft}$ is either (1) an indicator variable equal to one if the firm has workers in other unobserved divisions that are in the top quintile of the wage distribution and zero otherwise or (2) the fraction of the firm’s workers in unobserved divisions that are in the top quintile of the wage distribution. Our null hypothesis is that $\beta = 0$. That is, the presence of high-wage workers in other unrelated divisions of a firm does not affect the benefits that the firm provides to its other workers. Because $HiWageExcDiv_{ft}$ does not vary across different job hits from the same firm and because there is likely to be serial correlation in benefit costs, we cluster standard errors at the firm level in all regressions.

We report the results of estimating Equation 3 in Table 5. In Columns 1 to 3, we use the extensive margin measure of $HiWageExcDiv_{ft}$ (i.e., the indicator for the presence of workers

²⁰Importantly, we do not observe the exact mapping from wages to wage-related benefits defined by the plan, making it difficult to remove the expected variation in benefits based on wage levels, which need not be smooth or linear.

in excluded divisions within the firm that have wages in the top quintile of the distribution). In Columns 4 to 6, we use the fraction of workers in excluded divisions that are in the top quintile of the wage distribution, exploiting also intensive margin variation in the prevalence of high-wage workers inside the firm. Note that the fixed effects isolate contemporaneous comparisons of jobs from establishments in the same 3-digit NAICS industries and from establishments in the same state. In all specifications, we also include granular (6-digit) occupation fixed effects. These fixed effects are important to isolate as closely as we can variation in $HiWageExcDiv_{ft}$ between workers in jobs that are otherwise identical. For this reason, we also always include the natural logarithm of the average hourly wage paid to workers in the job.²¹ For example, the SOC code 43-9021 is for “data entry keyers.” Our regressions would then identify the differences in benefits paid to two data entry keyers in different firms, accounting for any differences in their hourly wages, but who have colleagues in other divisions of their respective firms with different wage profiles. Importantly, we also restrict the sample only to conglomerate firms that have excluded divisions operating in different industries from the establishments in our sample. This restriction ensures that our estimate of β will not pick up a conglomerate firm effect.

We find that the presence of high-wage workers in unrelated divisions of the firm indeed has a positive association with worker benefits, confirming Prediction 1. On the extensive margin (Column 1), the presence of high-wage workers elsewhere in the firm is associated with a roughly 11.5% increase in transformed health benefits from its mean value (1.701) or, alternatively, 26% of a standard deviation. The result is statistically significant at the 1% level.²² In Column 2, we add additional controls for the size of the firm and the establishment at which the job is located. We do not find any additional effect of firm size, conditional on inclusion in our sample of jobs from conglomerate firms; however, we do find that larger establishments pay higher health benefits. We also include a control for firm age. We find that older firms tend to pay higher benefits, though the effect is only marginally significant. The relation between

²¹In a small number of cases, the average hourly wage is imputed. We include these observations to avoid sample attrition. However, the wage information on which $HiWageExcDiv_{ft}$ is based is never imputed, nor are benefit costs.

²²The corresponding change in non-transformed benefits relative to its mean is roughly 23%.

benefits and $HiWageExcDiv_{ft}$ is qualitatively unchanged. In Columns 4 and 5, we report the results of the parallel estimations using the fraction of high-wage workers in excluded divisions to measure $HiWageExcDiv_{ft}$. The results are similar. In Column 4, a one standard deviation increase in the fraction of high-wage workers is associated with a roughly 3% increase in health benefit costs from their mean value, or 7% of a standard deviation.

Another channel that could generate between-firm variation in benefits (and wages) is positive assortative matching between workers and firms. In this mechanism, high-productivity workers sort into high-productivity firms and a firm effect in benefits (or wages) reflects the resulting differences in productivity. A challenge for this interpretation of our evidence from Section 4.1 is the consistently larger firm effect on benefits than on wages. It is unclear why firms would disproportionately compensate higher productivity using generous benefit packages, which can have different value to different employees and are difficult to individually tailor to the worker, rather than using wages.²³ In the context of Equation 3, high wages should at least partially reflect higher productivity of the workers in the job at that establishment. Thus, a starting point to control for sorting is to control directly for job-level wages.

We perform two additional tests to distinguish further the matching channel from the mechanism in Prediction 1. First, we introduce a direct control for firm-level labor productivity, measured as the natural logarithm of revenues divided by total employment. We report the results in Columns 3 and 6 of Table 5. We find some evidence of a positive association between labor productivity and health benefits (the relation is marginally significant in Column 3, but insignificant in Column 7). However, in both cases, we continue to find a strong positive association between the presence of high-wage employees elsewhere in the firm and health benefits, even conditional on labor productivity. It is also worth noting that controlling for labor productivity comes at the cost of a reduction in the usable sample because firm revenues are only available in the data for roughly 70% of firms.²⁴ Second, we report the results of

²³Note that the value of a generous health benefits package to an employee can also change dynamically and unpredictably given changes in the employees personal circumstances (e.g., due to the birth of a child, a personal health shock, etc.). Yet, firms do not dynamically adjust benefit packages in response to such changes so that the value of the worker's compensation package remains fixed.

²⁴Note that the differences in the estimated coefficients on $HiWageExcDiv$ in Columns 3 and 7 compared to

including a binary control for high-wage jobs (i.e., jobs with average hourly wages higher than the overall sample median) and its interaction with $HiWageExcDiv_{ft}$. If our estimate of β were confounded by the positive assortative matching mechanism, then the effect should be the strongest for the highest productivity workers (here, captured as the workers in the highest wage jobs) assuming that there is less variation in productivity among low-wage counterparts. Instead, we estimate a negative interaction effect in both regressions, though the estimates are not statistically significant (Columns 4 and 8). That is, the association between the presence of high-wage employees elsewhere in the firm and the generosity of benefit packages is not weaker for lower paid (low productivity) jobs. While challenging to reconcile with the matching story, the direction of the effect is consistent with our proposed mechanism. Benefit costs to low wage workers converge upwards towards the equilibrium benefits that must be paid to attract and retain the firm’s most productive workers. Thus, even though the channels are not mutually exclusive, the within-firm convergence channel appears to be important to explain the additional firm-level variation that we observe in benefits relative to wages.

4.2.2 Benefit Costs and Firm Location

In Section 4.2.1, we provide direct evidence for the within-firm convergence of benefits towards the benefits offered to the firm’s highest paid workers (Prediction 2). For this channel to generate the excess between firm variation in benefits (relative to wages) that we observe in Section 4.1, the level of benefits in the firm must also vary according to the market conditions faced by the firm’s most highly compensated workers (Prediction 3). In this Section, we take this prediction to the data.

We adapt Equation 3, replacing $HiWageExcDiv_{ft}$ with a measure of the equilibrium benefit costs to workers at the firm’s headquarters. Specifically, we define the variable $HQHealth_{fit}$ as the inverse hyperbolic sine of the leave-one-out average of health benefits in the MSA and industry in which the firm’s headquarters location operates. We identify headquarters locations

the corresponding regression columns that do not include the productivity controls are almost entirely driven by the restriction in sample (i.e., they are observed after excluding observations in which productivity cannot be measured even without including the control in the regressions).

using the LBD as the physical establishment with the highest payroll per employee among all of the firm’s establishments.²⁵ Because the construction of this variable does not require worker-level wage information from the LEHD data, we consider a larger regression sample than in Section 4.2.1.²⁶

In Table 6, we report the regression estimates. In Column 1, we find a positive and significant association between worker benefits and the average benefits paid by other firms to their workers in the firm’s headquarters location and industry. We include the same fixed effects and controls as in Column 1 of Table 5. So, again, the variation that identifies the estimate comes from the contemporaneous comparison of workers in the same job across firms with different labor market characteristics in their headquarters locations and industries, partialling out the effect of job-level wages as well as fixed effects for state and industry. Economically, a one standard deviation increase in the leave-one-out average health benefits in the headquarters location and industry is associated with an increase in transformed benefits of roughly 4.2% from its in-sample mean (1.418), or roughly 7% of a standard deviation.²⁷ As in Table 5, we add additional controls for firm and establishment size and firm age in Column 2, with little effect on the estimate of β . We also add an indicator for high-wage jobs and its interaction with $HQHealth_{fit}$ in Column 3, here finding essentially no effect of the interaction term.

The results suggest that market conditions relevant to the firm’s high-skilled workers are indeed correlated with the benefits provided throughout the firm. However, providing more generous benefits to other workers within the firm, if they do not face the same market conditions, is costly to the firm. Thus, there is a tradeoff. Firms should be more likely to be willing to bear this cost if a larger fraction of their (high-skill) workforce is located at headquarters. To bring this conjecture to the data, we allow the correlation between health benefits and the average benefits paid by other firms in the firm’s headquarters location and industry to vary with the fraction of the firm’s workers that are employed at the headquarters establishment.

²⁵Note that the headquarters establishment identified in this way need not be part of the NCS sample.

²⁶Here the constraint on sample size comes from the requirement that we observe other establishments in the industry and MSA of the firm’s headquarters. To limit this constraint, we measure industries for this purpose at the 2-digit NAICS level.

²⁷The corresponding change in non-transformed benefits relative to its mean is roughly 7%.

We report the results in Column 4. Consistent with the conjecture, we estimate a positive and significant interaction term: there is a higher correlation of benefits with average benefits in the headquarters market when a higher fraction of firm employees works at headquarters.

In Columns 5 to 8 of Table 6, we restrict the sample to jobs from establishments that are located in a different MSA from the firm’s headquarters establishment. This restriction ensures that the local labor market in which headquarters workers’ equilibrium benefits are determined is different from the market in which the benefits of the workers in the regression sample are determined. We then repeat the four specifications from Columns 1 to 4 on the reduced sample. We find little difference in the estimate of β in Columns 5 and 6, compared to the corresponding estimates in Columns 1 and 2. However, in Column 7, we now find a significantly larger estimate for β among jobs with below-median wages than in jobs with above-median wages, consistent with the results from Section 4.2.1. We also continue to find a stronger correlation between health benefits and average benefits in the headquarters market when more of the firm’s workers are employed at headquarters (Column 8). Overall, the results in Table 6 provide strong confirmation of Prediction 2, linking the tendency for firms to equalize benefit payments across workers with the between-firm variation in benefits that we measure in Section 4.1.

5 Implications for Labor Markets and Firm Value

In Section 4 we establish the presence of a robust firm component of benefits that is distinct from a firm component of wages. These cross-firm differences can have implications for the distribution of human capital across firms and, as a result, for firm efficiency and profitability. To begin to unpack these channels, we measure differences in patterns of worker turnover across firms. Then, we test the predictions of our economic mechanism for labor demand and labor supply at different parts of the wage distribution as well as for firm value.

5.1 Baseline Relation between Benefits and Worker Turnover

We use worker level data from the LEHD program to measure quarterly turnover for establishments in our NCS sample. As in Section 4.2.1, we must restrict the analysis to the subset of establishments that are located in the states that are part of our LEHD data snapshot. To construct our measure, we first identify all full-quarter workers in each establishment-quarter who are between the ages of 16 and 99.²⁸ To filter out part-time workers (consistent with our job filters in the NCS data), we drop workers earning less than \$15,000 in annualized wages and for whom the job in question is not the only observed job that satisfies our sample restrictions. We then measure the turnover rate as the average of the entry and exit rates in the establishment, where exit requires that the worker is not observed again in the firm in a job that satisfies our sample restrictions for at least four quarters and entry requires that the last such job in which we observe the worker is in a different firm.²⁹

Given these measures of establishment-level turnover, we estimate variations of the following linear regression:

$$Turnover_{efist} = \beta Benefit_{efist} + \mathbf{X}'_{jefist} \boldsymbol{\omega} + \delta_s + \alpha_j + \mu_i + \theta_t + \epsilon_{jefist} \quad (4)$$

Because different occupations are likely to have different turnover rates, we estimate the regression at the job level so that we can control for differences in the distribution of occupations across establishments with fixed effects α_j . We also continue to control for state (δ_s), industry (μ_i) and year-quarter (θ_t) fixed effects, as in Equation 3. \mathbf{X}'_{jefist} is again a vector of control variables, which includes firm and establishment employment, firm age, and job-level wages (all in log form). Within an establishment, variation in benefits across jobs could be determined at least in part by expected differences in turnover rates. To remove this source of endogeneity, we

²⁸“Full-quarter” means that the worker received non-zero earnings in the establishment in the prior and following quarters. This restriction to “book-ended” quarters is necessary to be able to properly compare quarterly wages across workers because otherwise low wages could reflect payment for a partial quarter of work.

²⁹Note that the distinction between establishment and firm in this discussion is important. Firms can consist of many establishments and we filter out exits from one establishment of the firm to another in our measure of turnover.

measure benefits as the average health benefit across jobs in the same establishment.³⁰ Because health benefits are sometimes zero, we use the inverse hyperbolic sine function to transform average benefits rather than the natural logarithm. Our null hypothesis is $\beta = 0$; that is, the average benefits paid to workers in the establishment do not affect the establishment’s turnover rate. The alternative from Prediction 4 is that more generous benefits lead to lower worker turnover (i.e., $\beta < 0$). As elsewhere in our analysis, we cluster errors at the firm level to account for serial correlation as well as dependence of benefits across jobs and establishments in the firm.

In Column 1 of Table 7, we report the baseline results from estimating Equation 4. We find the expected associations between turnover and our included controls: We observe higher turnover in larger firms and establishments, younger firms, and jobs that pay less generous wages. We also observe a significant negative relation between establishment-level health benefits and turnover. A one standard deviation increase in the (inverse hyperbolic sine of) average benefits (0.8058) is associated with a 15% decrease in worker turnover from its mean value (0.06341), or roughly 9.4% of a standard deviation.

Even though taking establishment-level averages can remove job-level sources of endogeneity in benefits, the OLS estimate of β can still be subject to endogeneity concerns. For example, high benefits could correlate with a positive establishment culture or other establishment-level amenities not reflected in wages that make it a generally desirable place to work and tend to lower turnover rates. To address these remaining concerns, we exploit the evidence from Section 4.2.2 to construct an instrumental variable strategy. Specifically, we instrument for $Benefit_{efist}$ using $HQHealth_{fit}$, the leave-one-out average benefits paid by firms in the industry and MSA in which the firm’s headquarters operates.

To provide valid identification, $HQHealth_{fit}$ must be orthogonal to the error term in Equation 4. By predicting benefits with other firm’s benefit choices, we break the direct link to potential omitted firm factors, such as culture. Likewise, it is difficult to argue for a reverse

³⁰We average benefits to the level of SEINs rather than SEIN units, which correspond to physical establishments in our data.

causality channel by which establishment turnover could cause pre-determined average benefits paid by other firms in the headquarters industry and location. The most obvious threat to the exclusion criterion is the possibility that the benefits paid by other firms directly affect turnover at the establishment. However, the predicted direction of this effect would be positive (i.e., higher benefits outside the firm would predict higher turnover as workers seek to attain positions in those more generous firms). If $HQHealth_{fit}$ provides valid exogenous variation in $Benefit_{efist}$, then we instead expect to observe a negative relation with turnover.

In Column 2 of Table 7, we present the estimates of the first stage regression of our IV strategy. The estimates confirm the relevance of the instrument. We find a positive relation between $HQHealth_{fit}$ and $Benefit_{efist}$ that is statistically significant at the 1% level. The Kleibergen-Paap F-statistic for the first stage regression is 29, suggesting that the instrument is strong. We report the second stage estimates in Column 3. We confirm a negative effect of instrumented establishment-level benefits on turnover, also reinforcing our confidence that the instrument satisfies the exclusion restriction. That is, the channel that identifies β is the upward pressure on establishment-level benefits because the firm has to pay its executives high benefits to match what they can attain as part of their outside options. We find that the estimated economic magnitude of the effect is larger in the IV regressions than in the OLS specification in Column 1. Here an increase in benefits of the same size that we considered in our discussion of Column 1 would decrease turnover by roughly 49% from its mean, or roughly 31% of a standard deviation. It is important to note that theory indeed predicts that endogeneity should dampen the estimate of β in our application. The reason is that firms that expect low worker turnover will tend to pay lower benefits to their workers in equilibrium because they have less need to induce their workers to stay with the firm. Thus, the larger magnitude of the estimates in our IV specification is not a source of concern.

As mentioned above, one potential way that the exclusion restriction could fail in Column 3 is if the average benefits paid by the firm's industry peers in its headquarter location directly affect establishment turnover rates. The negative estimate of β suggests that this channel does not dominate the partial correlation that drives our identification. Nevertheless, we consider

an alternative approach to address the concern directly. Building on the results from Section 4.2.2, we reduce the likelihood of a direct link between the labor market of the establishment and the market in which headquarters employees benefits are determined by requiring that the firm’s headquarters is in a different MSA from the establishment in our sample. We then re-estimate the baseline OLS and IV specifications from Columns 1 to 3 on this subsample of establishments. We report the results in Columns 4 to 6 of Table 7. We again observe a negative relation between establishment-level benefits and turnover in the OLS specification. On this subsample, the standard deviation of establishment-level health benefits is smaller (0.5721). Thus, despite the differences in the estimates of β compared to the corresponding regressions on the full sample, the economic magnitudes of the estimated effects are similar. A one standard deviation increase in average benefits is associated with a decrease in turnover equal to roughly 10% of a standard deviation (compared to 9.4% in Column 1). Turning to the first-stage regression for the IV specification, we continue to observe a positive and statistically significant association between the leave-one-out average benefits paid by firms in the industry and MSA in which the firm’s headquarters operates and establishment-level benefits (Column 5), even restricting the sample to establishments in different MSAs from the headquarters location. We also again estimate a negative effect of establishment-level benefits on turnover rates in the second stage (Column 6). It is worth noting that the reduced sample size reduces the power of the first stage. Here, the Kleibergen-Paap Wald statistic is roughly 5.7. The economic magnitude of the estimates is somewhat larger than what we observed in the full sample but not hugely so, mitigating concern over potential weakness of the instrument. Here, a one standard deviation increase in average benefits is associated with a decrease in turnover equal to roughly 61% of a standard deviation (compared to 31% in Column 3). Thus, the results suggest that a direct link between the average benefits among firms in the headquarters industry and location is unlikely to drive the effects of establishment benefits on turnover that we estimate using our IV approach on the full sample.

Our analysis in Section 4.2.2 also shows that the relation between establishment health benefits and leave-one-out average benefits paid by firms in the industry and MSA in which the

firm's headquarters operates is stronger when employment at the headquarters establishments makes up a larger fraction of total firm employment. This evidence suggests that our IV approach will have a stronger first stage if we instead restrict our attention to establishments of such firms. In the remaining columns of Table 7, we report the same OLS and IV specifications on the subsample of firms for which employment at headquarters makes up at least 15% of total firm employment. This restriction is likely to have the effect of dropping establishments of the very largest and/or most geographically dispersed firms from our sample. We find a similar association between turnover and establishment-level benefits to the one on the full sample (Column 7). In Column 8, we confirm that the first stage relation between leave-one-out average benefits paid by firms in the industry and MSA in which the firm's headquarters operates and establishment-level health is stronger economically and statistically than the corresponding relation in the full sample. Here, the Kleibergen-Paap Wald statistic is roughly 32. Reassuringly, the economic magnitude of the negative effect of establishment-level benefits on turnover in Column 9 is similar to Column 3 (a 39% decrease from the mean turnover rate, or 25% of a standard deviation, for a one standard deviation increase in benefits).

Overall, we consistently find a negative and economically meaningful relation with turnover both in the aggregate and using a variety of approaches to isolate exogenous variation in benefits.

5.2 Implications for the Within-Firm Distribution of Workers

Average health benefits paid by a firm reduce worker turnover. In this Section, we unpack the relation between health benefits and turnover, analyzing differences in the relation in the cross-section depending on worker type. We focus on differences in workers wages (or skill), partitioning workers into three wage categories and then re-measuring turnover, entry, and exit separately within each group of employees. In particular, we isolate workers with hourly wages less than \$10 per hour. These workers are the least likely to have scarce skills. Thus, it is unlikely that high firm-level benefits packages, conditional on wages, are designed by the firm for the purpose of retaining workers in this group. For the opposite reason, we also isolate

workers with hourly wages greater than \$30 per hour. The final group consists of all other workers in the sample. Because worker information comes from the LEHD data, in which wages are reported at a quarterly level, we scale these cutpoints up by 40 hours per week and 52 weeks per year and apply them to annualized wages.

To begin, we test the second part of Prediction 4. Specifically, our economic channel predicts that the firms with more generous benefits offerings will face increased labor supply in particular from low-wage (skill) workers for whom the more generous benefits are a larger component of total compensation. To do so, we estimate Equation 4 three times, using establishment turnover rates in each of the three wage bins, respectively, as the dependent variable. In all regressions, we instrument for establishment-level health benefits using the leave-one-out average benefits paid by firms in the industry and MSA in which the firm's headquarters operates, building on the evidence from Section 5.1. We restrict the sample to establishments of firms in which at least 15% of the workers are employed at the headquarters location. This restriction isolates the subset of firms that are relatively more likely to be providing high benefits to attract their top talent and, therefore, should maximize the power of our analysis. We also consider a consistent sample of firms across the regressions; i.e., we require that firms employ workers in each of the three wage bins to be included in the sample so that turnover rates within each bin are defined. This restriction eases the interpretation of cross-regression comparisons, since the relations between benefits and turnover rates among subsets of workers are estimated within the same set of firms.

We report estimates of the common first stage of the IV regressions in Column 1 of Table 8. The estimates are very similar to those we report in Column 8 of Table 7, demonstrating the robustness of our results to different samples. Again, there is a robust and significant positive relation between leave-one-out average benefits in the industry-location of headquarters and establishment-level benefits. The Kleibergen-Paap Wald statistic is 28.07.

In Columns 2 to 4 of Table 8, we report estimates of the second-stage regressions using turnover rates among employees in the three wage bins, separately, as the dependent variables. First, we note that the mean turnover rate itself declines monotonically as we move from the

lowest to the highest wage bin: average turnover is 12.7% among the lowest wage workers, but only 3.7% among the highest wage workers. We also find that the effect of instrumented establishment-level benefits on turnover declines monotonically as we move from the lowest to the highest wage bin. To compute economic magnitudes, we consider the same increment to establishment benefits that we used in Columns 7 to 9 of Table 7 (0.808; or a standard deviation on the associated sample). We find that an increase of this size in benefits leads to a decrease in turnover of 5.9 percentage points among the lowest wage workers, a decline in turnover of roughly 46.5% from its mean value, or 35.2% of a standard deviation. For workers in the middle wage bin, the estimated effect on turnover is very similar in magnitude to the estimated effect on the overall turnover rate from Table 7, though it is not statistically significant (possibly due to lower power in the smaller sample). Specifically, a 0.808 increase in establishment benefits leads to a decrease in turnover of 2.3 percentage points, which is a decline of 36% from the mean of turnover among workers in the wage bin and 23.8% of a standard deviation. Among high-wage workers, the estimated effect of establishment benefits on turnover is not only statistically insignificant, it is also economically small (a 14% decline from the mean for a 0.808 increase in benefits, or 6% of a standard deviation).

Recall that our prior results suggest that firms choose high levels of benefits because they are necessary to attract and retain the firm's highest skilled workers (Section 4.2.2), but that those benefits then tend to be equalized across workers in the firm (Section 4.2.1). Our results in the first four columns of Table 8 suggest that this latter pattern leads to an outsized decrease in the turnover rate of low-skilled workers. This pattern is consistent with a labor demand- or supply-side channel. From the supply side, the (equalized) firm-level benefit is likely to have larger economic relevance to low-wage workers because benefits with a given dollar value make up a larger percentage of total compensation. Thus, the benefit could reduce turnover rates through the channel of reduced exit rates. From the demand side, the (equalized) firm-level benefit makes low-wage workers more expensive to the firm, particularly to the extent that they are unable to reduce wages to offset more generous benefit payments (note that minimum wage constraints can bind among this set of workers). Thus, the benefit could reduce turnover rates

through the channel of reduced hiring (i.e., lower entry rates).

To sharpen our interpretation of the economic channels driving the turnover results in the cross-section and to test the labor demand channel from Prediction 5, we next distinguish between turnover that is driven by entry into and exit from the establishment. We construct a set of binary indicators that take the value one for quarters in which an establishment's entry rate of workers in a given wage bin exceeds its exit rate. We then repeat the regressions from Columns 2 to 4 of Table 8 using these binary indicators in place of within-wage-bin turnover rates as the dependent variables. Thus, the regressions can be interpreted as linear probability models. We report the results of the second stage regressions in Columns 5 to 7. Among low-wage workers, we find evidence consistent with a strong demand-side channel. The effect of establishment-level benefits on the likelihood that entry into the establishment exceeds exit from the establishment is negative and significant. A 0.808 increase in benefits leads to a roughly 10.5 percentage point drop in the likelihood that entry exceeds exit, relative to the mean likelihood of 29%. Thus, reduced turnover among low-wage workers appears to come, at least in large part, from reduced hiring of low-wage workers into the firm (relative to departures). Interestingly, at the top end of the wage distribution, we find evidence that the insignificant relation between benefits and turnover hides economically meaningful effects of benefits on entry and exit patterns. Here, the estimated effect of establishment-level benefits on the likelihood of entry exceeding exit is positive and significant.³¹ A .808 increase in benefits leads to an increase in likelihood of roughly 20 percentage points, relative to the mean likelihood of 19%. Thus, this evidence suggests that benefits are indeed an effective tool for the firm to retain highly skilled workers. High-wage workers become much more likely to enter than to leave the firm as benefits become more generous. The effect of benefits on the relative entry likelihood in the middle of the wage distribution is not statistically significant, but points in the direction of increased attractiveness of the firm to workers.

Overall, the results are consistent with a firm response to the turnover dynamics created

³¹To reconcile the results to Table 7, note that the increase in turnover from higher entry apparently is largely offset by a decrease in exit.

by the equalization of benefits across workers within the firm. Offering the generous benefits required to attract and retain the firm’s highest skilled workers has the side effect of increasing the cost of the firm’s low-skill workers (and likely also disproportionately reducing their incentive to leave the firm). In response, firms appear to reduce their hiring of low skilled workers, even while generous benefits appear to serve the purpose of increasing the attractiveness of the firm to higher skilled colleagues. To the degree that firms still require the functions that were performed by the low-skilled workers, our results suggest that the cost of providing benefits could motivate “domestic outsourcing,” automation, or other purposeful means of labor-replacement within the firm; that is, firms that are bound to offer generous nonwage compensation to their employees because they must do so to attract needed high-skill labor choose to reduce low-skill employment to avoid extending costly benefits to additional full-time employees. Increases in the cost of providing benefits over time, then, could be a driver not only of phenomena such as increased automation or domestic outsourcing, but also of increased specialization of firms within the human capital distribution (e.g., a decline in the attractiveness of the conglomerate organizational form).

5.3 Firm Value Implications

Thus far we have shown that there is a firm component of benefits that affects patterns of labor supply across and within firms. These patterns are consistent with the constraints in the amount of worker-level customization of benefit offerings imposed upon firms by various regulatory and institutional friction. Thus, as a final step, we investigate the value implications of these patterns. As in the previous subsections, we focus on health benefits.

An existing literature documents a positive correlation between firm profits and wages.³² While higher profits paired with higher wages does not guarantee higher firm value, evidence suggests we should expect a positive relationship. Michelacci and Quadrini (2009) argue that

³²Blanchflower, Oswald, and Sanfey (1996) look at industry wage differentials and find industries that pay higher wages also have higher profits per employee. Hildreth and Oswald (1997) show a positive firm-level correlation between profits and wages. More recently, Abowd, Kramarz, and Margolis (1999) use employee-employer matched data to document that high-wage firms are more profitable, on average.

financially constrained firms offer lower wages and numerous studies have shown that financially constrained firms have lower firm values. Perhaps, the most direct evidence comes from Mueller, Ouimet, and Simintzi (2017). In their paper, the authors find a positive correlation between Tobin’s Q and pay inequality, where higher wages for high-skill workers will increase pay inequality. Thus, to the extent that higher benefits reflects more skilled labor, we should expect to see a positive correlation between benefits and firm value.

However, nonwage benefits are distinct, vis-a-vis wages, due to the pressures to provide more equalized benefits across employees within a firm. This inflexibility can result in firms providing higher benefits to some employees than would otherwise be optimal. Firms may be able to offset some of these higher costs by offering lower wages, but minimum wage laws, downwards wage rigidities, and differences in the salience and relative valuation of nonwage benefits as well as the need for firms to provide market-clearing wages can limit this option. We have already shown evidence suggesting that firms mitigate these costs by reducing demand for low-wage workers. However, it may not be possible to replace all low-wage workers, in which case high-benefit firms may be over-compensating low-skill workers, leading to relatively lower valuations (Prediction 6).

To investigate value implications, we consider the firm’s scaled market value, measured by Tobin’s Q. Specifically, we estimate the following linear equation:

$$Q_{jfst} = \beta \text{Benefit}_{jfst} + \mathbf{X}'_{jfst} \boldsymbol{\omega} + \delta_s + \alpha_j + \mu_i + \theta_t + \epsilon_{jfst} \quad (5)$$

We measure Q as the log-transformed ratio of the market value of assets divided by the book value of assets. To control for differences in benefits across occupations as well as differences in firm value given occupational employment, we include occupation fixed effects α_j . We also control for industry (μ_i) times year-quarter (θ_t) and state (δ_s) times year-quarter (θ_t) fixed effects to control for time-varying patterns in Q . \mathbf{X}'_{jfst} is again a vector of control variables, which may include firm employment, firm age, and job-level wages (all in log form). Because health benefits are sometimes zero, we use the inverse hyperbolic sine function to transform

average benefits rather than the natural logarithm. As elsewhere in our analysis, we cluster errors at the firm level to account for serial correlation as well as the dependence of benefits across jobs and establishments in the firm.

We report the results in Table 9. In Column 1, we show a negative correlation between Q and health benefits. A one standard deviation increase in health care costs per employee is correlated with a decrease in Q of 6.8%, relative to the sample mean. Given the positive correlation between benefits and wages reported in Table 2, we next control for wages. Interestingly, controlling for wages does not attenuate the negative correlation between benefits and Q (Column 2). Moreover, we document a positive relationship between wages and Q , consistent with the earlier literature. We find similar results if we include controls for firm size and age (Column 3). While we are only able to show correlations, the combination of a positive correlation between Q and wages and a negative correlation between Q and benefits rules out many endogenous interpretations of our findings. For example, our results are not consistent with a mechanism through which firms with low Q pay excessively high *total* compensation, given the positive correlation between wages and Q . To further understand how the form of compensation matters, we consider a different independent variable, *BenefitRatio*, which is the ratio of nonwage benefits to total compensation. Consistent with the earlier specifications, we find a negative correlation between Q and the ratio of nonwage benefits to total compensation (Column 4).

Finally, our rigidity mechanism suggests that generous benefits should be most costly for firms employing relatively more low-wage workers that are hard to replace. In order to attract and retain high-wage employees, firms may be obligated to provide generous benefits, which they will then be required to also provide to their low-wage workers. To the extent that firms cannot easily replace low-wage workers with technology or through other means and instead provide benefit compensation in excess of what they would otherwise offer, this will put downward pressure on firm value. To date, manual and abstract jobs have been harder to replace with automation, while routine jobs have been easier to replace (Autor and Dorn (2013)). Moreover, abstract jobs tend to be high-skill, and hence high wage, while manual jobs tend to be lower

wage jobs and routine jobs tend to be more in the middle of the wage distribution. As such, we predict that the negative relation between benefits and Q will be most pronounced among firms with more manual workers and least pronounced among firms with more abstract workers.

We identify jobs as being manual, routine or abstract based on the tasks associated with each occupation as defined in Autor and Dorn (2013). Specifically, we code variables *manual*, *routine*, and *abstract* if the manual, routine or abstract tasks for that occupation are above the sample median, 0 otherwise.³³ Then, in column 5, we interact *BenefitRatio* with whether the occupation is manual, routine or abstract. We find a negative and significant coefficient on the interaction with manual. This result is consistent with these workers tending to have low skill, and hence low wages, but also being more difficult to replace. On the other hand, we find a positive coefficient on the interaction with abstract jobs, consistent with high-skill workers being in a labor market in which high benefits are required to attract and retain top talent.

In sum, we find a negative correlation between non-wage benefits and firm value. While these results are simple correlations, they are consistent with the earlier findings in the paper. Pressures to equalize nonwage benefits within the firm create a cost for the firm. While firms try to mitigate these costs by reducing their reliance on low-wage workers, such actions appear to be an imperfect fix, leaving high benefit firms with relatively lower valuations compared to their low benefit peers.

6 Conclusion

Nonwage benefits are an increasingly important part of total worker compensation, yet the extant literature has mostly ignored benefits due to lack of data. Using a novel database that matches job-level BLS data on benefit costs with Census employer-employee matched data, we document less variation in health and retirement benefits within firms than in wages. By contrast, between-firm variation in health, retirement and leave benefits is two to three times larger than between-firm variation in wages.

³³For occupations with missing values, we assign a value of 0.

To jointly explain these facts, we appeal to nondiscrimination regulations as well as fairness concerns and administrative costs associated with managing multiple plans, all of which exert pressure on firms to equalize benefits across workers in the firm. We then note that many firms operate in multiple labor markets. Employees in one market (segmented by skills or location) may command higher benefits than workers in another due to market pressures. Thus differences in the configuration of labor markets in which firms operate can generate a firm-specific component to the benefits they pay to workers. Consistent with this prediction, we show that the firms that employ high-bargaining power workers are also the firms that provide the most generous health benefits to their workers across the board.

This mechanism for determining benefits has important consequences for firms' low-wage employees. On the one hand, these employees receive more generous benefits than they would otherwise command in the marketplace. However, firms also readjust, leading to changes in turnover dynamics. Requirements to keep nonwage benefits more consistent within the firm make it costly to employ both high-wage and low-wage workers. We document evidence that firms with generous health benefits decrease their reliance on low-wage workers. Factors that lead firms to offer more standardized benefits within a firm advantage low-wage workers by making it more likely they will receive the same benefits as workers with more bargaining power, but also increase layoff risk and/or reduce employment opportunities. Moreover, they are likely to contribute to increasing firm specialization by skill levels. Thus, formally estimating the contribution of benefits to worker compensation can be essential to understanding patterns in the distribution of human capital across firms and over time as well as changing patterns in inequality.

References

- Abowd, J. M., F. Kramarz, and D. N. Margolis. 1999. High Wage Workers and High Wage Firms. *Econometrica* 67:251–333.
- Autor, D. H., and D. Dorn. 2013. The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market. *American Economic Review* 103:1553–97.
- Autor, D. H., F. Levy, and R. J. Murnane. 2003. The Skill Content of Recent Technological Change: An Empirical Exploration. *Quarterly Journal of Economics* 118:1279–333.
- Barth, E., A. Bryson, J. C. Davis, and R. Freeman. 2016. It’s Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States. *Journal of Labor Economics* 34:S67–97.
- Berk, J., R. Stanton, and J. Zechner. 2010. Human Capital, Bankruptcy and Capital Structure. *Journal of Finance* 65:891–926.
- Blanchflower, D., A. Oswald, and P. Sanfey. 1996. Wages, Profits, and Rent-Sharing. *Quarterly Journal of Economics* 111:227–51.
- Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler. 2023. How Much Should we Trust Estimates of Firm Effects and Worker Sorting. Working Paper.
- Brown, C., and J. Medoff. 1989. The Employer Size-Wage Effect. *Journal of Political Economy* 97:1027–59.
- Brown, C., and J. L. Medoff. 2003. Firm Age and Wages. *Journal of Labor Economics* 21:677–97.
- Buchmueller, T. C., and R. G. Valletta. 1996. The Effect of Health Insurance on Married Female Labor Supply. *Journal of Human Resources* 34:42–70.
- Bureau of Labor Statistics. 2021. *Employer Costs for Employee Compensation September 2021*.
- Cole, A., and B. Taska. 2023. Worker Valuation of Retirement Benefits. Working Paper, MIT.
- Dube, A., W. Lester, and M. Reich. 2016. Minimum wage shocks, employment flows, and labor market frictions. *Journal of Labor Economics* 34:663–704.
- Dube, A., S. Naidu, and A. Reich. 2022. Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers. Working Paper, University of Massachusetts.
- Ellul, A., M. Pagano, and F. Schivardi. 2018. Employment and Wage Insurance within Firms: Worldwide Evidence. *Review of Financial Studies* 31:1298–340.
- Freeman, R. B. 1981. The Effect of Unionism on Fringe Benefits. *Industrial and Labor Relations Review* 34:489–509.
- Gao, J., S. Ge, L. Schmidt, and C. Tello-Trillo. 2023. How Do Health Insurance Costs Affect Firm Labor Composition and Technology Investment? Working Paper, Georgetown University.

- Garcia, D., and O. Norli. 2012. Geographic Dispersion and Stock Returns. *Journal of Financial Economics* 106:547–65.
- Goldschmidt, D., and J. F. Schmieder. 2017. The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure. *Quarterly Journal of Economics* 132:1165–217.
- Graham, J. R., K. Kim, Hyunseob, S. Li, and J. Qiu. 2022. Employee Costs of Corporate Bankruptcy. Working Paper, Duke University.
- Haltiwanger, J., H. Hyatt, and J. Spletzer. 2022. Industries, Mega Firms, and Increasing Inequality. Working Paper, University of Maryland.
- Hamermesh, D. 1999. Changing Inequality in Markets for Workplace Amenities. *Quarterly Journal of Economics* 114:1085–123.
- Hau, H., Y. Huang, and G. Wang. 2020. Firm response to competitive shocks: Evidence from China’s minimum wage policy. *The Review of Economic Studies* 87:2639–71.
- Hildreth, A., and A. Oswald. 1997. Rent-Sharing and Wages: Evidence from Company and Establishment Panels. *Journal of Labor Economics* 15:318–37.
- Kristal, T., Y. Cohen, and E. Navot. 2020. Workplace Compensation Practices and the Rise in Benefit Inequality. *American Sociological Review* 85:271–97.
- Lamadon, T., M. Mogstad, and B. Setzler. 2022. Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review* 112:169–212.
- Liu, T., C. A. Makridis, P. Ouimet, and E. Simintzi. 2022. The Distribution of Non-Wage Benefits: Maternity Benefits and Gender Diversity. *Review of Financial Studies* forthcoming.
- Madrian, B. C. 1994. Employment-Based Health Insurance and Job Mobility: Is There Evidence of Job Lock? *Quarterly Journal of Economics* 109:27–54.
- Maestas, N., K. Mullen, D. Powell, T. von Wachter, and J. Wenger. 2023. The Value of Working Conditions in the United States and Implications for the Structure of Wages. *American Economic Review* forthcoming.
- Marinescu, I., Y. Qiu, and A. Sojourner. 2021. Wage Inequality and Labor Rights Violations. Working Paper, University of Pennsylvania.
- Mas, A. 2017. Does Transparency Lead to Pay compression? *Journal of Political Economy* 125:1683–721.
- Michelacci, C., and V. Quadrini. 2009. Financial Markets and Wages. *Review of Economic Studies* 76:795–827.
- Mitchell, O. S. 1982. Fringe Benefits and Labor Mobility. *Journal of Human Resources* 17:286–98.
- Mueller, H., P. Ouimet, and E. Simintzi. 2017. Within Firm Pay Inequality. *Review of Financial Studies* 30:3605–35.

- Oyer, P. 2008. Salary or Benefits? *Research in Labor Economics* 28:429–67.
- Silva, R. 2010. Internal Labor Markets, Wage Convergence and Investment. *Journal of Financial and Quantitative Analysis* 56:1192–227.
- Sockin, J. 2022. Show Me the Amenity: Are Higher-Paying Firms Better All Around? Working Paper, U.S. Treasury.
- Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter. 2019. Firming Up Inequality. *Quarterly Journal of Economics* 134:1–50.
- Taber, C., and R. Vejlin. 2020. Estimation of a Roy/Search/Compensating Differential Model of the Labor Market. *Econometrica* 88:1031–69.
- Tate, G., and L. Yang. 2015. The Bright Side of Corporate Diversification: Evidence from Internal Labor Markets. *Review of Financial Studies* 28:2203–49.
- Woodbury, S. 1983. Substitution between Wage and Nonwage Benefits. *American Economic Review* 73:166–82.

A Appendix

Health and retirement benefits are tax-advantaged, relative to cash compensation. For example, regular cash wages are typically tax deductible from the firm’s point of view (thereby lowering firm federal income taxes), but taxable from the employee’s point of view (leading to higher individual federal income taxes). Employer contributions to a 401(k) plan are typically tax deductible from the firm’s perspective and any individual tax implications can be deferred until distribution.³⁴ Likewise, premiums paid by employers for health insurance reduce firms’ tax liability, but do not create tax liabilities for individuals.

In order to qualify for these tax benefits, firms must meet non-discrimination standards, which act to ensure that tax-qualified nonwage benefits do not accrue disproportionately to the highly compensated individuals within the firm. Similar to the US tax code, these nondiscrimination tests are complicated. For example, a 401(k) sponsor must annually pass the Actual Deferral Percentage (ADP) test and the Actual Contribution Percentage (ACP) test. The ADP and ACP tests look for discrimination by comparing average deferrals and contributions of highly compensated employees (HCEs) to non-highly compensated employees (NCHEs). An HCE is defined as an employee who meets either of the following criteria: 1) owns more than 5% of the firm; or 2) earns more than \$135,000 a year and was in the top 20% of employees ranked by compensation.³⁵ Most plans must also pass the Top-Heavy test each year. A plan will fail the top-heavy test if the value of assets in accounts held by key employees sums to more than 60% of total plan assets, where “key employees” are defined as officers earning \$200,000 or more or company owners. To make things more complicated, firms can still benefit from preferential tax treatment, but avoid these tests, if they qualify as a “safe harbor 401(k).” Such plans are required to satisfy a mandatory contribution requirement either as a match or nonelective contribution. A common example would provide a 100% match on the first 3% of deferred wages and a 50% match on the next 2% of deferred wages. Such a plan would still lead to compression in benefits given high wage earners are subject to maximum deferral limits,

³⁴See <https://www.irs.gov/retirement-plans/plan-sponsor/401k-plan-overview> for more information.

³⁵\$135,000 is the 2022 cutoff. This threshold is revised each year.

currently indexed at \$20,500.

IRS nondiscrimination rules for health plans are similar but only apply to self-insured health plans and arrangements where employees contribute pre-tax dollars to pay for certain employee benefits, typically referred to as Section 125 or cafeteria plans. Employer-sponsored health insurance plans must also comply with other federal laws that ban discrimination, such as the Americans with Disabilities Act (ADA), the Age Discrimination in Employment Act (ADEA), and the Health Insurance Portability and Accountability Act (HIPAA), which all limit the ability to design health plans that have different eligibility requirements, contributions, or benefits across workers within a firm.

Leave benefits are unique as there is no specific nondiscrimination regulation that applies. However, there may still be pressures to equalize leave benefits within a firm to reduce the administrative burden of managing multiple plans or out of fairness concerns, which could be particularly acute for leave benefits given the visibility of days of absence.



Figure 1. The plots show the distributions of either wages or total compensation (measured as wages and health, retirement and leave benefits). Compensation is measured in \$/hour and normalized to \$2014.




Figure 2. The figures plot the between- and within-variance decomposition of nonwage benefits and wages by year. Between (Within) variance is the equal-weighted annual average of total between- (within-) establishment variance, estimated quarter-by-quarter. Between Variance for nonwage benefits is plotted using the left-hand side axis; Other measures are all plotted using the right-hand side axis. Nonwage benefits are inverse hyperbolic sine transformed and wages are log transformed. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages that can be matched to the LBD. We exclude the public sector (NAICS92). Panel A (health sample) drops observations in which health benefits are imputed. Panel B (retirement sample) drops observations in which retirement benefits are imputed. Panel C (leave sample) drops observations in which leave benefits are imputed.

Table 1. Summary Statistics. Means with standard deviations in parentheses. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages which can be matched to the LBD. We exclude the public sector (NAICS92). Column 1 drops observations where health benefits are imputed. Column 2 drops observations where retirement benefits are imputed. Column 3 drops observations where leave benefits are imputed. *Has Benefit* is an indicator variable taking the value of 1 if the benefit is available. *Wage + Benefits* is calculated as the sum of wages and health, retirement and leave benefits. All dollar values are normalized to 2014\$. Sample sizes are rounded to nearest 1,000.

	Health (1)	Retirement (2)	Leave (3)
<i>Job-level Characteristics</i>			
Has Benefit (%)	87.4 [33.1]	72.8 [44.5]	95.5 [20.7]
Benefit (\$/hour)	3.17 [2.38]	1.67 [3.54]	2.59 [2.71]
Wage (\$/hour)	26.03 [19.46]	25.48 [19.20]	25.22 [18.58]
Wage + Benefits (\$/hour)	33.90 [25.05]	32.94 [24.61]	32.61 [23.56]
Hours Worked / Year	1906 [234.8]	1906 [235.1]	1917 [231.4]
Unionized (%)	13.1 [33.8]	11.6 [32.0]	11.5 [31.9]
Employees / Job	28.51 [122.1]	26.51 [115.1]	25.41 [106.4]
<i>Establishment-level Characteristics</i>			
Establishment Employment	1149 [3015]	1075 [3038]	981 [2700]
Jobs Sampled / Establishment	4.71 [3.26]	4.60 [3.19]	4.46 [3.08]
<i>Firm-level Characteristics</i>			
Firm Employment	37050 [146900]	44730 [186200]	20190 [52900]
Firm Age	27.52 [8.56]	27.27 [8.73]	27.60 [8.70]
Jobs Sampled / Firm	24.24 [71.00]	26.93 [92.07]	14.67 [33.77]
Number of unique observations	519,000	545,000	523,000

Table 2. Pairwise Correlations. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages which can be matched to the LBD. We exclude the public sector (NAICS92). Column 1 drops observations where health benefits are imputed. Column 2 drops observations where retirement benefits are imputed. Column 3 drops observations where leave benefits are imputed. Benefits are inverse hyperbolic sine transformed. Wages are log transformed. All dollar values are normalized to 2014\$. *** denotes statistical significance at the 1% level.

	Health Benefit (\$/hour)	Retirement Benefit (\$/hour)	Leave Benefit (\$/hour)
	(1)	(2)	(3)
Wage (\$/hour)	0.469***	0.614***	0.746***
Unionized (%)	0.309***	0.276***	0.111***
Establishment Employment	0.416***	0.418***	0.366***
Firm Employment	0.356***	0.334***	0.333***
Firm Age	0.283***	0.220***	0.214***
Health Benefit (\$/hour)	-	0.527***	0.526***
Retirement Benefit (\$/hour)	0.541***	-	0.572***
Leave Benefit (\$/hour)	0.542***	0.574***	-

Table 3. Variance Decomposition. *Between Variance* is the time-series average of total between-firm variance, estimated quarter-by-quarter. *Within Variance* is the time-series average of total within-firm variance, estimated quarter-by-quarter. *Total Variance* is the sum of *Between Variance* and *Within Variance*. *% Between Variance* is the percent of *Total Variance* that is attributed to between-firm variance. Variance decomposition statistics apply to nonwage benefits (inverse hyperbolic sine transformed) in columns 1-4 and wages (log transformed) in columns 5-8. All statistics in a given row are estimated on the same sample. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages which can be matched to the LBD. We exclude the public sector (NAICS92). Row 5 drops observations where health benefits are imputed. Row 6 drops observations where retirement benefits are imputed. Row 9 drops observations where leave benefits are imputed. All dollar values are normalized to 2014\$.

	Benefits				Wages			
	Between Variance (1)	Within Variance (2)	Total Variance (3)	% Between Variance (4)	Between Variance (5)	Within Variance (6)	Total Variance (7)	% Between Variance (8)
Panel A. Health and Retirement								
Health	0.604	0.096	0.700	86.3%	0.184	0.128	0.312	59.0%
Retirement	0.694	0.113	0.807	86.0%	0.185	0.126	0.311	59.5%
Panel B. Leave								
Leave	0.426	0.159	0.585	72.8%	0.181	0.116	0.297	60.9%

Table 4. Variance decomposition after demeaning by occupation, state, or industry. *Between Variance* is the time-series average of total between-firm variance, estimated quarter-by-quarter. *Within Variance* is the time-series average of total within-firm variance, estimated quarter-by-quarter. *Total Variance* is the sum of *Between Variance* and *Within Variance*. *% Between Variance* is the percent of *Total Variance* that is attributed to between-firm variance. Variance decomposition statistics apply to nonwage benefits (inverse hyperbolic sine transformed) in columns 1-4 and wages (log transformed) in columns 5-8. All statistics in a given row are estimated on the same sample. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages which can be matched to the LBD. We exclude the public sector (NAICS92). Rows 5-7 drops observations where health benefits are imputed. Rows 10-12 drop observations where retirement benefits are imputed. Rows 15 and 17 drop observations where leave benefits are imputed. In rows 5, 10 and 15 we estimate variance decomposition statistics after demeaning the observations by state. In rows 6, 11, and 16 we estimate variance decomposition statistics after demeaning by occupation (6-digit SIC code). In rows 7, 12 and 17 we estimate variance decomposition statistics after demeaning by industry (6-digit NAICS). All dollar values are normalized to 2014\$.

	Benefits				Wages			
	Between Variance	Within Variance	Total Variance	% Between Variance	Between Variance	Within Variance	Total Variance	% Between Variance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Health Benefits								
...Demean State	0.565	0.100	0.665	85.0%	0.164	0.128	0.292	56.2%
...Demean Occupation	0.397	0.134	0.531	74.8%	0.054	0.049	0.103	52.4%
...Demean Industry	0.281	0.099	0.380	73.9%	0.068	0.126	0.194	35.1%
Panel B. Retirement Benefits								
...Demean State	0.653	0.115	0.768	85.0%	0.165	0.126	0.291	56.7%
...Demean Occupation	0.403	0.125	0.528	76.3%	0.053	0.048	0.101	52.5%
...Demean Industry	0.294	0.116	0.410	71.7%	0.070	0.124	0.194	36.1%
Panel C. Leave Benefits								
...Demean State	0.388	0.159	0.547	70.9%	0.160	0.116	0.276	58.0%
...Demean Occupation	0.175	0.109	0.284	61.6%	0.055	0.046	0.101	54.5%
...Demean Industry	0.158	0.158	0.316	50.0%	0.072	0.114	0.186	38.7%

Table 5. Relation between Health Benefits and High Wage Workers. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed health benefits which can be matched to the LBD. We exclude the public sector (NAICS92). We also require that all of the firm's establishments are observed in our snapshot of the LEHD. The dependent variable is the dollar value of health benefits (inverse hyperbolic sine transformed). In columns 1-4, *HiWageExcDiv* is estimated as an indicator variable which takes the value of 1 if the firm has workers in other unobserved divisions that are in the top quintile of the wage distribution. In columns 5-8, *HiWage* is defined as the fraction of the firm's workers in unobserved divisions there in the top quintile of the wage distribution. *Labor Productivity* is measured as the natural logarithm of firm revenues divided by firm employment. *Job_AboveMed* is an indicator variable which equals 1 if the given job is in the top half of the wage distribution. Wages, firm and establishment employment, and firm age are log transformed. Standard errors are clustered at the firm level. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.10$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HiWageExcDiv	.1973*** [0.0451]	.1765*** [0.0414]	.1222** [0.0482]	.1963*** [0.0434]	.1846*** [0.0674]	.1603** [0.0665]	.2216*** [0.0768]	.1648** [0.0775]
Wage	.5173*** [0.0546]	.4816*** [0.0528]	.4902*** [0.0646]	.4311*** [0.0566]	.522*** [0.0543]	.4829*** [0.0521]	.4798*** [0.0629]	.4321*** [0.0562]
Firm Employment		0.00587 [0.0132]	.03743* [0.0201]	0.006349 [0.0132]		0.01624 [0.0139]	.0477** [0.0203]	0.0165 [0.0138]
Firm Age		.07813* [0.0441]	.1636*** [0.0613]	.07826* [0.0439]		0.06621 [0.0449]	.1596*** [0.0614]	0.0661 [0.0447]
Establishment Employment		.0583*** [0.0150]	0.01806 [0.0186]	.05803*** [0.0149]		.05752*** [0.0153]	0.0142 [0.0188]	.05741*** [0.0153]
Labor Productivity			.04276* [0.0257]				0.03437 [0.0254]	
Job_AboveMed				.1201** [0.0577]				.08669** [0.0406]
HiWage * Job_AboveMed				-0.04602 [0.0597]				-0.01049 [0.0871]
NAICS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occ FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs	41000	41000	28500	41000	41000	41000	28500	41000
Adj. R2	0.5037	0.5131	0.5354	0.514	0.4986	0.5094	0.5363	0.5102

Table 6. Relation between Health Benefits and Local Compensation. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-imputed health benefits which can be matched to the LBD. We exclude the public sector (NAICS92). We also require that all of the firm's establishments are observed in our snapshot of the LEHD. In columns 5-8, we further limit the sample by requiring that the establishment is located in a different MSA as the firm's headquarters. The dependent variable is the dollar value of health benefits (inverse hyperbolic sine transformed). *HQHealth* is the leave-one-out average of the health benefits in the MSA and industry in which the firm's headquarters location operates (inverse hyperbolic sine transformed). *Job_AboveMed* is an indicator variable which equals 1 if the given job is in the top half of the wage distribution. *Frac Emp at HQ* is the fraction of total firm employment that works at the headquarters location. Wages, firm and establishment employment, and firm age are log transformed. Standard errors are clustered at the firm level. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.10$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
HQHealth	.1005*** [0.0143]	.08419*** [0.0131]	.0832*** [0.0159]	0.004061 [0.0211]	.09603*** [0.0344]	.09829*** [0.0352]	.1837*** [0.0491]	.08543** [0.0331]
Wage	.6661*** [0.0278]	.5488*** [0.0249]	.4748*** [0.0274]	.5474*** [0.025]	.1717*** [0.0375]	.1694*** [0.0371]	.1280*** [0.0329]	.1663*** [0.0363]
Firm Employment		.02674*** [0.0054]	.02672*** [0.0054]	.02289*** [0.0062]		-0.01528 [0.0132]	-0.01622 [0.0129]	-0.009373 [0.0139]
Firm Age		.1511*** [0.0171]	.1505*** [0.017]	.1496*** [0.017]		-0.1077 [0.1498]	-0.1131 [0.1469]	-0.08435 [0.1465]
Establishment Employment		.06546*** [0.0069]	.06587*** [0.0069]	.06631*** [0.007]		0.01168 [0.0077]	0.0112 [0.0076]	0.0102 [0.0072]
Job_AboveMed			.1166*** [0.0374]				.3072*** [0.0971]	
HQHealth * Job_AboveMed			0.0001665 [0.0198]				-.1111** [0.0494]	
Frac Emp at HQ				-.237*** [0.0471]				-1.063** [0.4586]
HQHealth * Frac Emp at HQ				.1196*** [0.0254]				.6784*** [0.2541]
NAICS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occ FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs	297000	297000	297000	297000	40500	40500	40500	40500
Adj. R2	0.4026	0.4409	0.4424	0.4423	0.6765	0.6787	0.6823	0.6818

Table 7. Relation between Health Benefits and Turnover Rates. The unit of observation is at the establishment job quarter level. Establishment-level data is aggregated across full-quarter workers earning \$15,000 or more annually in establishments in our LEHD snapshot and matched to NCS. The dependent variable is the quarterly establishment-level *turnover rate* (columns 1 and 3; 4 and 6; 7 and 9) and *Estab_Health* (columns 2, 5, and 8). Turnover is measured as the average of the entry and exit rates. *Estab_Health* is the average health benefit at the establishment. *HQHealth* is the leave-one-out average of the health benefits in the MSA and industry in which the firm's headquarters location operates (inverse hyperbolic sine transformed). Columns 1 and 4 are OLS. Columns 2, 5, and 8 report first-stages of IV2SLS regressions. Columns 3, 6, and 9 report second-stage regressions. Wages, firm and establishment employment, and firm age are log transformed. Standard errors are clustered at the firm level. (***) p<0.01, (**) p<0.05, (*) p<0.10).

	<u>Full Sample</u>			<u>HQ MSA ≠ Estab MSA</u>			<u>HQ Emp ≥15% Firm Emp</u>		
	OLS	IV2SLS First-stage	Second- stage	OLS	IV2SLS First-stage	IV2SLS Second-stage	OLS	IV2SLS First-stage	Second- stage
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Wage	-.005133*** [0.0016]	.5072*** [0.0329]	0.008775 [0.0083]	0.00225 [0.0026]	.09816*** [0.0308]	.0129** [0.0064]	-.005876*** [0.0018]	.5356*** [0.0316]	0.005313 [0.0088]
Estab_Health	-.01179*** [0.0012]		-.03858** [0.016]	-.01794*** [0.0053]		-.1248** [0.0587]	-.01137*** [0.0013]		-.03171** [0.0158]
HQHealth		.09937*** [0.0185]			.08106** [0.034]			.1088*** [0.0192]	
Firm Employment	.001337*** [0.0004]	.01993*** [0.0063]	.001936*** [0.0006]	0.002374 [0.0015]	-.03112** [0.0155]	-0.0003371 [0.0026]	.001811*** [0.0006]	.03772*** [0.009]	.002604*** [0.0009]
Firm Age	-.01269*** [0.002]	.1794*** [0.0236]	-.007828** [0.0034]	-0.02128 [0.0221]	-.3682* [0.2216]	-0.06602 [0.0406]	-.01474*** [0.0021]	.19*** [0.0247]	-.01084*** [0.0036]
Establishment Employment	.001936*** [0.0007]	.07272*** [0.0091]	.00384*** [0.0014]	-0.001502 [0.001]	0.0144 [0.009]	-0.0000799 [0.0011]	.002398** [0.0009]	.05734*** [0.0117]	.003552*** [0.0013]
NAICS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occ FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald		28.96			5.69			32.12	
Obs	137000	137000	137000	16500	16500	16500	117000	117000	117000
Adj. R2	0.4190	0.5031		0.7666	0.7873		0.4076	0.4861	

Table 8. Relation between Health Benefits, Turnover and the Relative Rate of Entry and Exit by Wage Bin. The unit of observation is at the establishment quarter wage-bin level. Regressions in Columns 2 through 7 are the second stages of IV2SLS specifications. Column 1 reports the (common) first stage regression. Columns 2 and 5 report results for the lowest wage bin. Columns 3 and 6 report results for the middle wage bin. Columns 4 and 7 report results for the highest wage bin. The dependent variables are *Estab_Health* (column 1), *Turnover* (columns 2-4) and *More Entry than Exit* (columns 5-7). *Turnover* is measured for the establishment-quarter as the average of the entry and exit rates of workers into and out of the establishment. *More Entry than Exit* is a dummy variable equal to one if the entry rate of workers into the establishment (in the relevant wage bin) exceeds the exit rate of workers from the establishment. *Estab_Health* is the average health benefit at the establishment. Wages, firm and establishment employment, and firm age are log transformed. Standard errors are clustered at the firm level. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.10$).

	<u>First Stage</u>	<u>Second Stage</u>					
		<i>Turnover Rate</i>			<i>More Entry than Exit</i>		
		Wage Bin 1	Wage Bin 2	Wage Bin 3	Wage Bin 1	Wage Bin 2	Wage Bin 3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Wage	.4322*** [0.0326]	.03606** [0.0143]	0.007648 [0.0088]	0.002151 [0.0068]	0.02207 [0.036]	-0.05163 [0.0456]	-0.04301 [0.0455]
Estab_Health		-.07336** [0.0321]	-0.02832 [0.0198]	-0.006607 [0.0151]	-.1296* [0.0755]	0.1416 [0.0993]	.2496** [0.0973]
HQHealth	.1212*** [0.0229]						
Firm Employment	.03234*** [0.0113]	.003868** [0.0018]	0.001371 [0.001]	0.0009302 [0.0008]	0.00485 [0.0045]	0.000883 [0.0053]	-0.001409 [0.0056]
Firm Age	.175*** [0.0294]	-0.008496 [0.007]	-.009671** [0.0042]	-.007722** [0.0032]	-0.02503 [0.0163]	-.04406** [0.0199]	-.06715*** [0.0207]
Establishment Employment	.03843*** [0.0146]	.01044*** [0.0025]	.003235** [0.0014]	0.0008634 [0.0012]	.05875*** [0.006]	.04354*** [0.007]	.03762*** [0.0067]
NAICS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occ FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap Wald	28.07						
Obs	74000	74000	74000	74000	74000	74000	74000
Adj. R2	0.5014						

Table 9. Relation between Health Benefits and Firm Value. The unit of observation is at the establishment quarter wage-bin level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed health benefits which can be matched to the LBD. The dependent variable is Q, measured as the log transformed ratio of the market value of assets divided by the book value of assets. Health is the estimated cost of providing health benefits, inverse hyperbolic sine transformed. BenefitRatio is the ratio of nonwage benefits to total compensation (wages and nonwage benefits). Routine, Manual and Abstract are indicator variables if the tasks typically performed by the occupation are in the top 1/2 of the sample distribution in terms of being routine, manual or abstract, respectively, as described in Autor and Dorn (2013). Wages, firm and establishment employment, and firm age are log transformed. Standard errors are clustered at the firm level. (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

	(1)	(2)	(3)	(4)	(5)
Health	-.03978*	-.04281**	-.04185**		
	[0.0212]	[0.021]	[0.0208]		
Wage		.02222***	.02325***		
		[0.0084]	[0.0083]		
Firm Employment			-0.008626	-.008927*	-.00892*
			[0.0053]	[0.0053]	[0.0053]
Firm Age			-0.04996	-0.05233	-0.05205
			[0.0635]	[0.064]	[0.064]
BenefitRatio				-.3747***	-.4344***
				[0.1375]	[0.1652]
BenefitRatio * Routine					0.1084
					[0.076]
BenefitRatio * Manual					-.1697*
					[0.095]
BenefitRatio * Abstract					.2147*
					[0.1217]
NAICS3 * Year-Quarter FE	Yes	Yes	Yes	Yes	Yes
State * Year-Quarter FE	Yes	Yes	Yes	Yes	Yes
Occ FE	Yes	Yes	Yes	Yes	Yes
Obs	121000	121000	121000	121000	121000
Adj. R2	0.8194	0.8196	0.8205	0.8206	0.8208

Internet Appendix Table 1. Establishment-level Variance Decomposition. *Between Variance* is the equal-weighted time-series average of total between-establishment variance, estimated quarter-by-quarter. *Within Variance* is the time-series average of total within-establishment variance, estimated quarter-by-quarter. *Total Variance* is the sum of *Between Variance* and *Within Variance*. *% Between Variance* is the percent of *Total Variance* that is attributed to between-establishment variance. Variance decomposition statistics apply to nonwage benefits (inverse hyperbolic sine transformed) in columns 1-4 and wages (log transformed) in columns 5-8. All statistics in a given row are estimated on the same sample. The unit of observation is at the establishment job quarter level. The sample includes all observations in the NCS between 2004-2014 for full-time employees with non-missing data and non-imputed wages which can be matched to the LBD. We exclude the public sector (NAICS92). Row 6 drops observations where health benefits are imputed. Row 7 drop observations where retirement benefits are imputed. Row 8 drop observations where leave benefits are imputed. All dollar values are normalized to \$2014.

	Benefits				Wages			
	Between Variance (1)	Within Variance (2)	Total Variance (3)	% Between Variance (4)	Between Variance (5)	Within Variance (6)	Total Variance (7)	% Between Variance (8)
Health	0.619	0.096	0.715	86.6%	0.183	0.133	0.316	57.9%
Retirement	0.757	0.126	0.883	85.7%	0.185	0.132	0.317	58.4%
Leave	0.428	0.16	0.588	72.8%	0.184	0.118	0.302	60.9%