NBER WORKING PAPER SERIES

WHY ARE SUPERSTAR FIRMS THE BEST PLACES TO WORK? COMPETITIVE ADVANTAGE IN AMENITY PROVISION AND THE DECLINING LABOR SHARE

Ioannis Branikas
Briana Chang
Harrison Hong
Nan Li

Working Paper 33870
http://www.nber.org/papers/w33870

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2025

We thank Philip Bond, Yihui Pan, Chendi Zhang, Youchang Wu, Ron Giammarino, Lorenzo Garlappi, Jessica Jeffers, Andrew Chamberlain, Brian Blank, Jincheng Tong, Andrianos Tsekrekos, Elisa Ughetto, Andreas Milidonis, Andreas Charitou, Gabriel Buchbinder, Daniel Zhao, and seminar participants at UT Dallas, UNSW, University of Oregon, Mississippi State University, University of Idaho, University of Cyprus, AFA, Pacific Northwest Finance Conference, and Exeter Sustainable Finance Conference for many helpful comments. We are also grateful to Glassdoor.com for providing data. Branikas acknowledges support from the Marilyn C. and Gerry B. Cameron Center for Finance and Securities Analysis. Li acknowledges support from the Social Sciences and Humanities Research Council of Canada. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Why Are Superstar Firms the Best Places to Work? Competitive Advantage in Amenity Provision and the Declining Labor Share
Ioannis Branikas, Briana Chang, Harrison Hong, and Nan Li
NBER Working Paper No. 33870
May 2025
JEL No. G0, G3, G39, J3, J31, J33

ABSTRACT

Why are superstar firms the best places to work? We document that worker satisfaction scales nearly twice as fast as wages with firm size. Our assignment model distinguishes selection (talent sorting) from competitive advantage (amenity provision); amenity-cost advantages cause utility to rise more steeply than wages. Using Glassdoor data, competitive advantage explains roughly half of the utility increase with firm size. As workers increasingly value non-pecuniary amenities, superstar firms with a cost advantage substitute wages for amenities, raising profits and reducing the measured labor share all while raising worker utility. Declining labor shares need not signal deteriorating worker welfare.

Ioannis Branikas
University of Oregon
Department of Finance
branikas@uoregon.edu

Briana Chang
Hong Kong University of Science
and Technology (HKUST)
Department of Finance
brianachang@ust.hk

Harrison Hong
Columbia University
Department of Economics
and NBER
hh2679@columbia.edu

Nan Li
University of Toronto
Nanjulie.Li@Rotman.Utoronto.Ca

1. Introduction

Large, dominant "superstar firms"—defined by their high productivity and an increasing concentration of market share (Autor et al. 2020)—consistently lead the rankings of the world's best places to work (see, e.g., Fortune or Glassdoor surveys). Companies such as Nvidia, Google, and Apple represent a new class of market leaders that combine operational excellence with superior human capital management. While a long-standing literature in labor economics documents a significant large-firm wage premium (Brown and Medoff 1989, Abowd et al. 1999, Mueller et al. 2017), this anecdotal evidence suggests that these superstar firms might offer an even more substantial amenity premium. Employees at these market leaders rate them highly on non-pecuniary dimensions such as work-life balance, corporate culture, and career development. Prior research confirms that these high-satisfaction workplaces are not merely incidental; they are strongly correlated with superior stock returns and long-term firm value (Edmans 2011, Edmans et al. 2024), suggesting that a firm's "superstar" status is fundamentally linked to a critical intangible: its workplace culture.

We begin by documenting a striking stylized fact that characterizes the modern labor market: worker satisfaction scales nearly twice as fast with firm size as cash wages do. To measure the relationship between worker satisfaction, we merge *Glassdoor* ratings with firm financials. Following Sockin 2024 and Sockin and Sojourner 2023, we interpret Glassdoor average ratings by workers of their firms as a monotonic transformation of workers' utility. We find that while the elasticity of wages with respect to firm assets is 0.045, the elasticity of the overall workplace rating is 0.079. This "slope differential" presents a central puzzle: why do superstar firms possess such an advantage in providing the best workplaces?

We distinguish between two primary mechanisms to explain this premium. The first is competitive advantage: superstar firms may possess a structural organizational technology or scale efficiency that allows them to supply costly amenities—such as flexibility, prestige, or professional development—more efficiently than smaller rivals. The second is selection: even in the absence of a cost advantage, productive firms attract high-talent workers. If these workers possess non-pecuniary preferences, the observed "best workplace" status is a byproduct of the sorting process—the firm is simply compensating its superstars in a currency they value. Distinguishing these channels is essential for understanding whether superior workplace culture is an active driver of firm value or a passive reflection of workforce composition.

To address this question, we develop an assignment model with endogenous amenities. Sorting between firms and workers is determined by standard production complementarities (Rosen 1981, Sattinger 1993, Chiappori and Salanié 2016) and by firm-level heterogeneity in the cost of providing amenities. In our framework, firm scale, wages, amenities, and worker utilities emerge endogenously as firms compete for talent. Importantly, a cost advantage for large firms causes utility to rise more steeply than wages, consistent with the stylized fact that workers' utility rises much faster with size than does wages. Hence, our model delivers a sharp empirical test: variation in wages isolates the portion of utility driven by worker selection, while any residual growth in utility—the "amenity premium"—identifies the firm's competitive advantage in amenity provision.

We can then use this differential-slopes test to back out underlying structural parameters that allow us to apportion the contribution of selection versus competitive advantage. Quantitatively, the structural competitive advantage of large firms explains approximately 60% of the variation in worker utility with firm size, while selection accounts for the remaining 40%. These results suggest that, while superstar firms certainly attract workers who demand better culture, these firms are also fundamentally more efficient at producing that culture.

Crucially, we analyze the implications of this advantage for firm profitability—a strategic dual to the worker-utility view. We show that when large firms possess a cost advantage in supplying amenities, they can effectively substitute expensive cash wages for efficiently produced non-pecuniary rewards. We estimate that a 1% increase in workers' non-pecuniary preferences corresponds to a 0.6% increase in firm profits. These findings identify workplace culture as a critical intangible that drives the superstar firm advantage.

A large and influential literature attributes the secular decline in the labor share to rising firm market power (Autor et al. 2020). As superstar firms capture increasing market share, the prevailing narrative holds that a shrinking fraction of value added is returned to workers — implying that the labor share decline is a direct signal of deteriorating worker welfare. This market power channel is well-documented and we do not dispute it. Our findings, however, point to a complementary mechanism with strikingly different welfare implications. As workers increasingly demand non-pecuniary amenities (Mas and Pallais 2017, He et al. 2021) — a trend accelerated by Covid-19 — firms with a scale-based cost advantage in amenity provision optimally substitute away from cash wages toward amenities. This substitution mechanically reduces the pecuniary labor share recorded in national accounts, generating the same reduced-form pattern as the market power story, but through an entirely different channel. We make no attempt to separately identify the two mechanisms empirically; rather, our point is that rising amenity demand alone is capable of producing a declining labor share, and that when it does, the welfare implications are more nuanced.

Specifically, higher amenity demand has three effects on worker welfare operating simultaneously. It increases the bargaining power of firms with a cost advantage in amenity provision, shifting surplus toward firms and lowering workers' share of output (profit-sharing effect). But it also raises the value of matching with a superstar firm, increasing the total surplus available to be shared (productivity effect), and fattens the upper tail of the firm-size distribution, disproportionately benefiting workers matched to the largest firms (fat-tail effect). Whether the average worker is better or worse off depends on which forces dominate — and critically, on a condition on firm scale that must be satisfied for the productivity and fat-tail effects to prevail. This condition is not guaranteed in general, but we show in Section 7 that it holds given our estimated parameters, implying that the average worker is better off even as the labor share declines.

Related literature and contribution. By identifying workplace amenities as a scalable technology for rent capture in superstar firms, this paper contributes to four major literatures: (i) assignment models, (ii) labor and finance, (iii) the labor share, and (iv) corporate culture.

We extend talent assignment models for labor and finance markets (Gabaix and Landier 2008, Tervio 2008, Sørensen 2007, Chang and Hong 2019, Pan 2017) by endogenizing the supply of amenities. While these models focus on wage scaling, our framework explains why worker satisfaction scales nearly twice as fast as wages ($\epsilon_U \approx 0.079$ vs $\epsilon_W \approx 0.045$). Our structural decomposition finds that a firm-side competitive advantage in amenity provision (κ_n) explains 60% of this premium, while worker selection accounts for the remaining 40%.

We formalize amenities as a scalable technology that allows firms to substitute expensive wages for efficiently produced amenities. This provides a micro-foundation for the finding by Edmans 2011 that employee satisfaction predicts superior stock returns; superstar firms capture rent by providing utility more efficiently than the market prices it. Our estimated 0.6% profit elasticity with respect to amenity preference is consistent with "value-enhancing" view of workplace quality.

Our paper contributes to the literature on labor and finance (Whited, 2019) by emphasizing non-wage compensation (Liu et al. 2023, Ouimet and Tate 2023) as a determinant of competition for worker talent. While compensating differentials for amenities are well documented (Rosen 1986, Lavetti 2023), the relative roles of selection and competitive advantage in shaping worker utility—and the implications for firm profits—have not been analyzed.[1]

Moreover, we address the consequences of the secular decline of the labor share documented by Autor et al. 2020 for worker utility (see Grossman and Oberfield 2022 for a review). The literature argues that this declining labor share leads to lower worker utility due to great firm power. When there are scale effects to the provision of workplace amenities,

[1]See also Hwang et al. 1998, Hwang et al. 1992, Sorkin 2018, Maestas et al. 2023 for empirical identification of compensating differentials.

declining labor shares need not presage lower worker utility when there is increasing demand for workplace amenities.

More broadly, our framework complements work that treats corporate culture as an intangible asset with direct value implications (Guiso et al. 2015, Li et al. 2021) and the broader human-capital-as-capital perspective (Eisfeldt et al. 2023). Our paper also has implications for the growing literature on corporate social responsibility and labor markets, where corporate mitigation can also be viewed as an amenity (Chang and Hong 2026, Colonnelli et al. 2024).

2. Data

The universe of firms in our study consists of U.S.-headquartered, publicly traded firms that were members of the S&P 500 Index for at least one year between 2006 and 2023. We use two types of datasets. The first consists of data from Glassdoor on employees' wages and workplace satisfaction ratings for their firms. The second is firm financial data from Compustat.

2.1. Firm wages

Our data on firm wages are based on salary reports extracted from Glassdoor, a prominent online labor market platform where employees can anonymously disclose reviews about their firms and salaries. A salary review contains information about the name of the employer firm, the type of employment, the job title, the base pay, and the bonuses. An employee's total pay is readily obtained by adding the last two items together.

We consider only full-time employees working at those companies who report their job titles. The latter are classified into 1,245 occupation categories based on Glassdoor's proprietary machine learning algorithm. If users declare that they are former employees of the firm they rate, we assume that they report their pay from the last year of their employment. For current employees, we attribute the reported salary to the year when the review is provided. Our sample spans from the year 2006 to the end of 2022.

To capture the salaries of talented employees, we drop reviews with reported salaries below $30,000, which was approximately the median annual wage in the U.S. at the start of our sample period in 2006 and represents the bottom 10% of the original data distribution. In order to remove outliers, we also omit reviews that are in the top 1% of the original data distribution, corresponding to reported salaries higher than $500,000. These filters lead us to a sample of 2,477,126 salary reviews.

To obtain an estimate for the wage of a typical talented employee at a given firm in a specific year, we implement a multi-step aggregation process. Although Glassdoor reviews are considered representative at the industry level (e.g., Karabarbounis and Pinto, 2018, Dehaan et al., 2023), they are voluntary and less frequently updated at the firm level, which may cause certain occupations to be over- or under-represented. To address this, we calculate firm-level salaries based on the set of representative occupations identified for each industry.

Specifically, we first measure the relative frequency of occupations within each industry according to the two-digit Global Industry Classification Standard (GICS). In particular, we calculate the ratio of the number of reviews that each occupation in a given industry receives during the entire sample period to the total number of reviews in that industry. Occupations that make up less than 1% are dropped, since there is only sparse information about them and they are likely unimportant for that industry. The relative frequencies of the selected top occupations in each industry are subsequently re-scaled, so that they add up to one. In that way, every industry ends up being represented by its most descriptive occupations.

For example, as displayed in Table 1, the total number of reviews in the Information Technology sector is 498,741. Software engineers fill out most of these reviews (i.e., 72,385 or approximately 14.5% of the total). Consultants, other types of engineers (e.g., systems engineers, design engineers, application engineers, process engineers, and hardware engineers), managers (e.g., project managers, product managers, account managers, and directors), analysts (e.g., systems analysts and business analysts), and sales representatives each contribute about 1% to 3% of the total. Overall, there are 22 job titles that contribute more than 1% of the reviews in this industry. All together, these top job titles account for approximately 50% of the industry's total reviews. More broadly, we report in Table 2 the top 5 occupations for all 11 GIC two-digit industries. Each industry has around 20 top occupations, representing 37% to 52% of their total salary reviews.

We then calculate the average salary for each industry's selected top occupation in a given firm-year. If an occupation within a specific firm-year pair has more than three reviews, we use the average salary for that occupation in that firm-year. Whereas if an occupation within a specific firm-year pair has fewer than three (or no) reviews (but is still considered to be descriptive of the firm's industry, since it belongs to the industry's selected top occupations), we impute its salary from the industry-year average salary for that occupation.

Lastly, we derive the wage of the representative talented employee in a given firm-year by computing the weighted average of the firm's occupations' average salaries in that year, with weights equal to the frequencies of the selected top occupations in the firm's industry. All in all, our aggregation accounts for the fact that the importance (and therefore the wage) of different occupations varies across industries. For instance, firms in the Information

Table 1
The selected top occupations in the Information Technology industry

This table depicts the selected top occupations in the Information Technology industry.

(1) Rank	(2) Title of Occupation	(3) Number of Reviews	(4) Relative Frequency (%)	(5) Cumulative Relative Frequency (%)	(6) Selected	(7) Re-scaled Relative Frequency (%)
1	software engineer	72,385	14.5	14.5	✓	28.9
2	consultant	15,896	3.2	17.7	✓	6.4
3	engineer	14,316	2.9	20.6	✓	5.7
4	program manager	13,376	2.7	23.3	✓	5.3
5	manager	11,847	2.4	25.7	✓	4.7
6	project manager	11,123	2.2	27.9	✓	4.4
7	software developer	9,328	1.9	29.8	✓	3.7
8	product manager	8,529	1.7	31.5	✓	3.4
9	software engineer in test	7,794	1.6	33.1	✓	3.1
10	systems analyst	7,776	1.6	34.7	✓	3.1
11	systems engineer	7,564	1.5	36.2	✓	3
12	account executive	7,230	1.4	37.6	✓	2.9
13	business analyst	7,157	1.4	39	✓	2.9
14	sales representative	6,865	1.4	40.4	✓	2.7
15	design engineer	6,860	1.4	41.8	✓	2.7
16	applications engineer	6,836	1.4	43.2	✓	2.7
17	account manager	6,680	1.3	44.5	✓	2.7
18	management consultant	6,367	1.3	45.8	✓	2.5
19	analyst	5,898	1.2	47	✓	2.4
20	process engineer	5,782	1.2	48.2	✓	2.3
21	hardware engineer	5,334	1.1	49.3	✓	2.1
22	director	5,196	1	50.3	✓	2.1
23	associate	4,659	0.9	51.2		
24	financial analyst	4,482	0.9	52.1		
25	marketing manager	4,419	0.9	53		
26	enterprise architect	4,292	0.9	53.9		
27	sales engineer	3,884	0.8	54.7		
28	technical support engineer	3,857	0.8	55.5		
29	customer service representative	3,370	0.7	56.2		
30	business development associate	3,123	0.6	56.8		
31	programmer analyst	3,121	0.6	57.4		
32	product marketing manager	2,919	0.6	58		
33	sales manager	2,591	0.5	58.5		
34	data scientist	2,574	0.5	59		
35	business development manager	2,527	0.5	59.5		
...	...	...	...	...		
	Total	498,741	100	-		100

Table 2

Top 5 occupations in each two-digit GICS industry

This table reports the five most frequent occupations in each of the 11 two-digit GICS industries, based on the universe of Glassdoor salary reviews from full-time employees with reported job titles. For each industry, the header row lists the total number of reviews and the number of occupations selected (those with relative frequency $\geq$ 1%). For each occupation, we report its re-scaled relative frequency (the share of the occupation among the industry's selected top occupations); these re-scaled frequencies sum to 100% across all selected occupations within an industry. See Section 2.1 for details on the aggregation methodology.

Rank	Occupation	%	Rank	Occupation	%
Information Technology			*Communication Services*		
498,741 reviews; 22 occupations selected			200,449 reviews; 20 occupations selected		
1	software engineer	28.9	1	software engineer	26.6
2	consultant	6.4	2	sales representative	14.2
3	engineer	5.7	3	customer service rep.	7.3
4	program manager	5.3	4	account executive	4.7
5	manager	4.7	5	retail sales associate	4.3
Financials			*Energy*		
425,411 reviews; 21 occupations selected			44,546 reviews; 20 occupations selected		
1	software engineer	9.3	1	engineer	14.9
2	analyst	8.8	2	field service engineer	10.3
3	senior leader	8.3	3	geologist	5.8
4	customer service rep.	7.7	4	mechanical engineer	5.6
5	personal banker	7.5	5	project manager	5.5
Consumer Discretionary			*Materials*		
422,050 reviews; 18 occupations selected			38,932 reviews; 20 occupations selected		
1	software engineer	15.7	1	store manager	11.4
2	store manager	13.9	2	sales representative	10.1
3	warehouse worker	8.4	3	engineer	7.4
4	customer service rep.	7.1	4	assistant manager	7.2
5	manager	6.3	5	manager	6.5
Industrials			*Real Estate*		
324,272 reviews; 21 occupations selected			29,654 reviews; 25 occupations selected		
1	software engineer	12.9	1	leasing agent	10.6
2	engineer	9.4	2	property manager	8.1
3	systems engineer	8.0	3	project manager	7.1
4	sales representative	7.2	4	real estate manager	5.2
5	manager	5.2	5	maintenance technician	5.1
Health Care			*Utilities*		
263,924 reviews; 19 occupations selected			22,496 reviews; 15 occupations selected		
1	research scientist	12.2	1	engineer	24.0
2	software engineer	7.5	2	manager	9.3
3	pharmacist	7.2	3	analyst	8.5
4	customer service rep.	6.9	4	customer service rep.	8.2
5	manager	6.5	5	project manager	8.0
Consumer Staples					
206,651 reviews; 23 occupations selected					
1	store manager	13.7			
2	pharmacist	8.8			
3	pharmacy technician	7.3			
4	manager	7.2			
5	warehouse worker	6.1			

Technology industry hire more and pay more software engineers, as do firms in the Consumer Staples industry for store managers.

The summary statistics for the obtained firm wages are presented in the first row of Panel A of Table 3. On average, the representative talented employee received approximately $89K. The standard deviation was around $19K, while the median was slightly over $85K.

2.2. Workers' ratings of workplace satisfaction

In addition to writing salary reviews on Glassdoor, employees are asked to provide an overall rating of their employer, as well as separate ratings for compensation and benefits, and work-life balance of their firm. We use these ratings, which range from one to five stars, to assess the relative importance of wage and work-life balance amenities to workers at their firms. Following Sockin (2024), we interpret an employee's overall rating of her firm as a noisy proxy for the total utility she derives from wages and other non-pecuniary amenities. The reliability of Glassdoor as a workplace data source has been examined in detail by Sockin and Sojourner 2023. Related, papers (Green et al. 2019, Huang et al. 2020) have established that Glassdoor review ratings are informative about company fundamentals. Hence, while we might not observe the level of amenities supplied by a firm (in contrast to wages), the workers' workplace ratings nonetheless serve as an indirect measure.

We again restrict our sample to full-time employees who report a job title in their reviews. This yields a sample of 931,274 rating reviews during the period 2012–2022. We then estimate the representative talented employee's overall rating, compensation and benefits rating, and work-life balance rating at the firm-year level, in the same manner that we obtain an estimated wage. That is, the aggregation process of Section 2.1 is repeated to identify the top occupations in each industry. Subsequently, for each firm in each year, we calculate the weighted average of its occupation-level overall ratings, using again the frequencies of the selected top occupations in the firm's industry as weights. We follow the same procedure for employees' ratings of their firms' compensation and benefits, as well as work-life balance.

The summary statistics of the derived variables are also presented in Panel A of Table 3. The average overall rating assigned by the representative talented employee to her firm was slightly over 3.3 stars, as was the median. The average and median ratings for compensation and benefits were around 3.5 stars, whereas the average and median work-life balance ratings were both approximately 3.2 stars.

Table 3
Summary statistics

This table summarizes the annually observed variables in our sample. Panel A refers to firms' wage and amenities' star ratings from Glassdoor. $Wage$ is the wage of a firm's representative employee. $TotRating$ is the overall rating of a firm's representative employee. CB is the compensation and benefit rating of a firm's representative employee. WL is the work-life balance rating of a firm's representative employee. Panel B refers to firms' employee numbers, assets, and profits from Compustat. EMP is a firm's employee number. $Assets$ is a firm's assets. IB is a firm's income before extraordinary items. NI is a firm's net income. IB/EMP, and NI/EMP are the corresponding measures of a firm's profits per worker. The sample is an unbalanced panel of 730 U.S.-based firms that were members of the S&P 500 Index during the period 2006-2022.

	Mean	S.D.	Median	P5	P95
Panel A: Firms' salary and amenities' star ratings from Glassdoor					
$Wage$ ($)	88,972	19,252	85,301	65,799	127,164
$TotRating$	3.36	0.27	3.34	2.96	3.84
CB	3.50	0.26	3.48	3.10	3.92
WL	3.18	0.31	3.18	2.72	3.72
Panel B: Firms' number of employees, assets, and profits from Compustat					
EMP	48,233	122,443	17,483	1,800	191,000
$Assets$ (million $)	67,759	259,925	13,182	1,659	223,432
IB (million $)	1,678	5,163	603	-618	7,602
NI (million $)	1,694	5,180	611	-645	7,722
IB/EMP	63,821	442,256	33,557	-50,278	323,713
NI/EMP	66,099	449,409	34,083	-52,182	336,699

2.3. Firm assets and profits

We draw annual data on firms' assets, number of employees, income before extraordinary items, and net income from Compustat. In particular, income before extraordinary items and net income constitute two alternative ways of measuring a firm's profit. Since the estimation of workers' non-pecuniary preferences that we present below requires data on firms' profits per worker, we divide each of these variables by the number of employees. The corresponding summary statistics are presented in Panel B of Table 3. Ultimately, our dataset is an unbalanced panel of 730 distinct firms spanning from the year 2006 to 2022, and containing in total 8,534 firm-year observations.

3. Stylized Facts

We document how wages and worker satisfaction scale with firm size. Our main finding — that satisfaction scales nearly twice as fast as wages with firm size — constitutes the central empirical result that motivates our model.

3.1. Workers' Wages with Firm Scale

We estimate the following regression model:

$$\ln(Wage_{i,t}) = b_1^{Wage} \ln(Assets_{i,t}) + b_2^{Wage} \ln(IndustryMedianAssets_{i,t}) + \eta_i^{Wage} + \epsilon_{i,t}^{Wage}, \tag{1}$$

where $\ln(Wage_{it})$ is the natural log of firm i's wage in year t, $\ln(Assets_{i,t})$ is the natural log of firm i's assets in year t, $\ln(IndustryMedianAssets_{i,t})$ is the natural log of the median firm's assets in firm i's industry in year t, η_i^{Wage} is firm i's fixed effect, and $\epsilon_{i,t}^{Wage}$ is the error. This regression specification naturally arises from assignment models Gabaix and Landier 2008 to pick up both scale effects, captured by the median firm's assets, and sorting due to complementarities.

We present three specifications that differ in how they control for unobserved heterogeneity. Column 1 pools all variation — both within-firm over time and across firms — without any fixed effects. Column 2 adds industry fixed effects, absorbing time-invariant differences in worker skill composition across industries. Column 3 is our preferred specification, which includes firm fixed effects. This is crucial: firm fixed effects isolate within-firm variation over time, allowing us to identify how wages and worker satisfaction scale with firm size as firms grow, rather than simply comparing large and small firms cross-sectionally. The latter comparison is confounded by composition effects — larger firms may systematically attract different types of workers — which could bias the estimated relationship between firm size and compensation in either direction. As we discuss below, this concern is not merely theoretical: the cross-sectional estimates tell a materially different story from the within-firm estimates, underscoring the importance of our preferred specification.

The results are presented in Panel A of Table 4. In all columns, the coefficient of $\ln(Assets_{i,t})$ is positive and statistically significant. However, the magnitude varies substantially across specifications in an economically meaningful way. Without fixed effects (Column 1), the estimated elasticity of wages with respect to firm assets is only 0.014. Once we control for firm fixed effects (Column 3), this estimate rises to 0.045 — more than three times larger — with a t-statistic of 5.63. The estimated coefficient of $\ln(IndustryMedianAssets_{i,t})$ is 0.100 with a t-statistic of 5.88, confirming that industry-level scale effects also play an important role. The fact that within-firm wage growth is substantially larger than the cross-sectional wage differential suggests that composition effects attenuate the cross-sectional estimates: as firms grow, they raise wages more than a simple comparison of large and small firms would suggest.

3.2. Worker Satisfaction with Firm Scale

We next study how workplace ratings scale with firm size using an analogous regression:

$$\ln(Rating_{i,t}) = b_1^{Rating} \ln(Assets_{i,t}) + b_2^{Rating} \ln(IndustryMedianAssets_{i,t}) + \ \eta_i^{TotRating} + \epsilon_{i,t}^{TotRating}, \tag{2}$$

where $\ln(Rating_{i,t})$ denotes the natural logarithm of either firm i's compensation and benefits rating ($\ln(CB_{i,t})$), overall rating ($\ln(TotRating_{i,t})$), or work-life balance rating ($\ln(WL_{i,t})$) in year t. Because we directly observe worker wages, the compensation and benefits rating is redundant; however, we report it for robustness. Following Sockin (2024), we focus on workers' overall (total) rating as our main measure of worker utility, since it reflects satisfaction with both wages and non-pecuniary amenities and is therefore the most directly applicable to our model. The work-life balance rating serves as an additional window into the amenity component specifically.

The same pattern of firm fixed effects mattering emerges — and more strikingly so. Panel B presents results for the compensation and benefits rating. In Column 3, the estimated coefficient of $\ln(Assets_{i,t})$ is 0.048 with a t-statistic of 8, broadly similar in magnitude to the wage elasticity of 0.045, consistent with the idea that the compensation and benefits rating largely reflects the pecuniary component of utility.

The central finding of the paper emerges in Panel C, which reports results for the overall

Table 4

Regressions of a firm's wage and workplace ratings on its own assets and the median firm's assets in its industry

This table presents the regressions of a firm's wage and workplace ratings on its own assets and the median firm's assets in its industry. In Panel A, the dependent variable is $\ln(Wage_{i,t})$, i.e., the log of firm i's wage in year t. In Panel B, the dependent variable is $\ln(CB_{i,t})$, i.e., the log of firm i's compensation and benefits rating in year t. In Panel C, the dependent variable is $\ln(TotRating_{i,t})$, i.e., the log of firm i's overall rating in year t. In Panel D, the dependent variable is $\ln(WL_{i,t})$, i.e., the log of firm i's work-life balance rating in year t. In all panels, the independent variables are $\ln(Assets_{i,t})$, i.e., the log of firm i's assets in year t, and $\ln(IndustryMedianAssets_{i,t})$, i.e., the log of the median firm's assets in firm i's industry in year t. The industries are the 11 two-digit GICS industries. The table reports the coefficient estimates and the two-way clustered standard errors at the firm and year level in parenthesis. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. See Table 3 for a detailed description of the sample.

	(1)	(2)	(3)
Panel A: Regressions of firms' wage			
$\ln(Assets)$	0.014**	0.007***	0.045***
	(0.005)	(0.002)	(0.008)
$\ln(IndustryMedianAssets)$	0.008	0.137***	0.100***
	(0.009)	(0.017)	(0.017)
Observations	8,534	8,534	8,534
R^2	0.016	0.235	0.332
Panel B: Regressions of firms' compensation and benefits rating			
$\ln(Assets)$	0.008***	0.007***	0.048***
	(0.002)	(0.002)	(0.006)
$\ln(IndustryMedianAssets)$	0.006	0.066***	0.043***
	(0.004)	(0.012)	(0.009)
Observations	4,713	4,713	4,713
R^2	0.038	0.093	0.224
Panel C: Regressions of firms' overall rating			
$\ln(Assets)$	0.007***	0.006***	0.079***
	(0.002)	(0.001)	(0.008)
$\ln(IndustryMedianAssets)$	-0.000	0.111***	0.071***
	(0.004)	(0.024)	(0.017)
Observations	4,747	4,747	4,747
R^2	0.015	0.128	0.298
Panel D: Regressions of firms' work-life balance rating			
$\ln(Assets)$	0.004*	0.003**	0.056***
	(0.003)	(0.001)	(0.005)
$\ln(IndustryMedianAssets)$	0.024***	0.076***	0.048***
	(0.005)	(0.016)	(0.011)
Observations	4,698	4,698	4,698
R^2	0.062	0.080	0.226
Panels A to D			
Industry FE	No	Yes	No
Firm FE	No	No	Yes

rating. Without firm fixed effects (Column 1), the estimated elasticity of overall ratings with respect to firm assets is only 0.007 — smaller than the corresponding wage elasticity of 0.014, suggesting that satisfaction rises *less* than wages with firm size. This pattern reverses sharply once we move to our preferred specification with firm fixed effects (Column 3): the elasticity of overall ratings rises to 0.079 with a t-statistic of 9.9, nearly *twice* the corresponding wage elasticity of 0.045. This reversal is telling. In the cross-section, larger firms attract workers whose observed satisfaction partially reflects their own characteristics rather than the firm's amenity provision. Once we control for firm fixed effects and isolate within-firm variation, the true scaling of worker utility with firm size is revealed to be substantially faster than wages — the slope differential that is the central puzzle of this paper.

Panel D reports results for work-life balance ratings. In Column 3, the estimated coefficient of $\ln(Assets_{i,t})$ is 0.056 with a t-statistic of 11.2, a value that lies between the wage elasticity and the overall rating elasticity. This is consistent with work-life balance capturing part of the amenity premium — more so than cash compensation, but less than the full utility gain reflected in the overall rating.

Graphical Evidence. We now present our regression evidence graphically. To facilitate clear visualization, within each two-digit GICS industry and year, we sort firms into three asset size groups — low, median, and high. For each industry-group-year, we calculate the log median wage and the log median asset size. To capture within-group variation, we regress each of these two variables on industry-group fixed effects and obtain the residuals. This approach is analogous to controlling for firm fixed effects in firm-level regressions.

In Subfigure 1a, we plot the residualized log wage against residualized log assets. The scatter plot shows a clear positive relationship, indicating that larger firms tend to pay higher wages within their industry group. In Subfigure 1b, the residualized compensation and benefits rating shows a similarly sloped fitted line, consistent with this rating reflecting primarily the pecuniary component of compensation.

The key contrast emerges in Subfigure 1c: the fitted line for the overall rating is visibly steeper than for wages, confirming graphically the slope differential documented in the regressions. Subfigure 1d shows that work-life balance ratings follow a similar pattern — steeper than wages but not as steep as the overall rating — consistent with work-life balance capturing part but not all of the amenity premium.

Takeaways. Two findings stand out. First, both wages and worker satisfaction increase with firm size, consistent with a large literature on the firm size premium. Second, and more importantly, the within-firm elasticity of worker satisfaction with respect to firm assets

Fig. 1. Scatterplots of firm wages and ratings against firm assets, after binning firm observations into three groups by industry–year. Each subfigure plots the relationship between a residualized firm-level workplace variable (log wage or rating) against the residualized log asset size. Firms are grouped into three asset bins within each two-digit GICS industry and year. Residuals are obtained by regressing the calculated log median wage, compensation and benefits rating, total rating, work-life balance rating and asset size on industry-group fixed effects. Subfigure 1a shows the residualized log wage. Subfigure 1b shows the residualized compensation and benefits rating. Subfigure 1c shows the residualized total rating. Subfigure 1d shows the residualized work-life balance rating.

(a) Residualized wage against residualized assets

(b) Residualized compensation and benefits against residualized assets

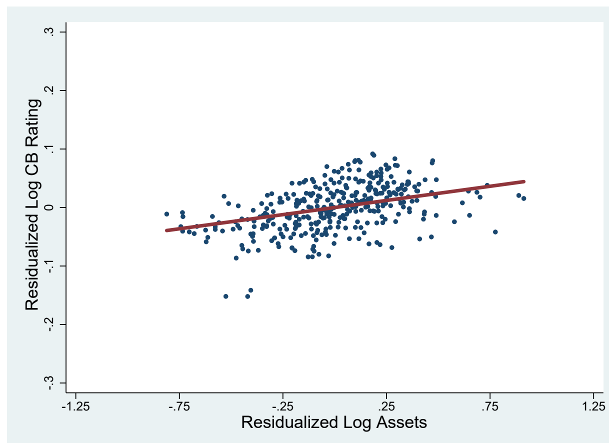

(c) Residualized overall rating against residualized assets

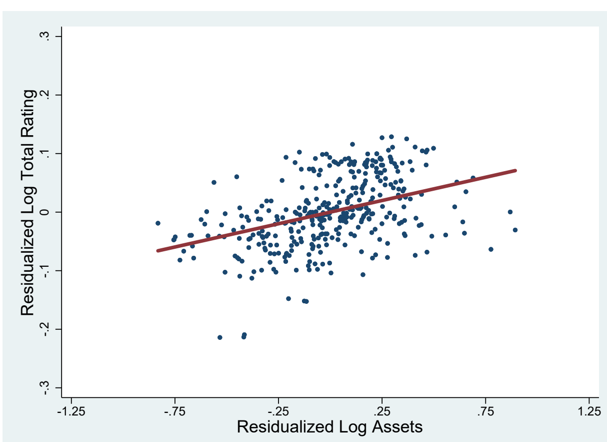

(d) Residualized work-life balance rating against residualized assets

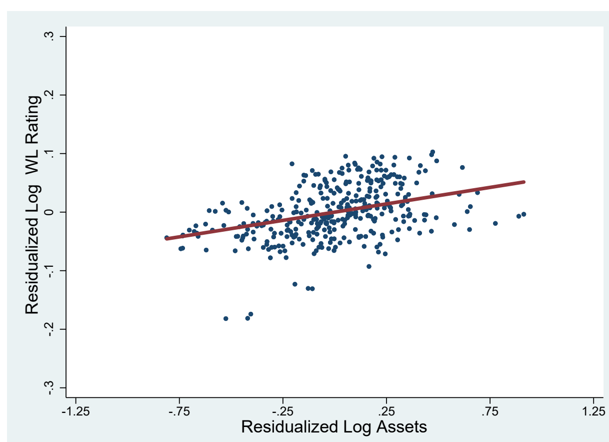

(0.079) is nearly twice that of wages (0.045). This slope differential is not visible in the cross-section — if anything, the cross-sectional estimates suggest the opposite — and only emerges once we isolate within-firm variation using firm fixed effects. This finding is the central empirical puzzle we seek to explain: why do superstar firms provide such a disproportionate utility premium relative to their wage premium, and what does this imply for the distribution of surplus between firms and workers?

4. Model

Section 3 documented a central puzzle: within-firm, worker satisfaction scales nearly twice as fast as wages with firm size. A standard assignment model — where larger firms attract more talented workers who are compensated purely in wages — cannot generate this slope differential. If utility and wages were proportional, they would scale identically with firm size. To explain why they diverge, we need a model in which firms compete on a dimension beyond wages, and in which larger firms have a structural advantage in that dimension.

We build such a model by extending the standard assignment framework in two ways. First, we allow workers to derive utility directly from non-pecuniary amenities — such as flexibility, workplace culture, and career development — in addition to wages. Second, and crucially, we endogenize the supply of these amenities by firms, allowing the marginal cost of provision to vary with firm size. This cost heterogeneity is the key ingredient: when larger firms can provide amenities more cheaply per worker, they optimally supply more of them, causing worker utility to rise faster than wages with firm size. The slope differential between satisfaction and wages therefore becomes an empirical signal of this competitive advantage — one we can use to separately identify the contributions of talent sorting (selection) and cost efficiency (competitive advantage) to the utility premium at superstar firms.

4.1. Environment

We consider an environment where heterogeneous firms and workers are matched in a competitive equilibrium each period. The parameters defined below are assumed to be time-invariant, except for the one referring to the productivity or size of the lowest-ranked firm in the economy, which acts as firm scale. Moreover, we assume that the matching is frictionless in every period, so that the matching decisions are effectively static. Based on this and for ease of notation, we omit the period subscript t in what follows.

Firms. There is a continuum of firms of mass 1. We use the terms of 'firm productivity' and 'firm size' (measured by assets) interchangeably. Firms are heterogeneous in their productivity or size. In particular, we assume that the firm size distribution is Pareto, with a tail index $\pi > 0$. Hence, the productivity of a firm with ranking k is given by $a[k] = a_L (1-k)^{-\frac{1}{\pi}}$. The time-varying scale of firms can be captured by changes in the parameter a_L over time. In other words, the equilibrium outcome should be interpreted given the value that a_L takes in a given period.

Workers. There is also a continuum of workers of mass 1. Workers are indexed by their skill s, with $F_w(s)$ denoting the measure of workers with skills below s. The distribution of worker talent is time invariant, and its ranking in the upper tail satisfies $s'[i] = B(1-i)^{-\beta-1}$. As discussed by Gabaix and Landier (2008), there exist constants β and B such that this expression holds. Indeed, we can define:

$$s[i] = \begin{cases} s_L(1-i)^{-\beta} & \text{if } \beta > 0 \\ s_H - \left(\dfrac{1-i}{\widetilde{B}}\right)^{-\beta} & \text{if } \beta < 0 \end{cases}. \tag{3}$$

The magnitude of β, which can be either positive or negative, captures the heterogeneity in talent. If $\beta > 0$, the worker talent distribution is Pareto with a tail index $1/\beta$. On the other hand, if $\beta < 0$, the distribution has an upper bound, denoted by s_H. As we show in Section 5, the data favor $\beta < 0$, implying that worker talent has an upper bound — a feature that will matter for the welfare analysis in Section 7.

Non-pecuniary preferences. The key departure from a standard assignment model is that workers care about more than wages. To capture this, we assume that utility is a Cobb-Douglas function of wages (denoted by x_0) and N different amenities (whose respective quantities are in the list $\{x_n\}_{n=1}^{N}$). Putting all the arguments in a vector $\boldsymbol{x}$, we have that:

$$u(\boldsymbol{x}) = \prod_{n=0}^{N} x_n^{\alpha_n}. \tag{4}$$

For instance, x_1 can represent the number of flexible hours, and x_2 the number of hours of remote work. The parameter $\alpha_n \in [0,1]$ captures the relative importance of amenity n in workers' utility, with $\sum_{n=0}^{N} \alpha_n = 1$. The weight on wages, α_0, and the total weight on amenities, $1-\alpha_0$, will be central to the labor share analysis in Section 7: as α_0 falls and workers place more value on amenities, firms optimally substitute away from cash wages, mechanically reducing the measured labor share.

Job positions. Depending on its productivity or size (a), a firm may have a different number of job positions to fill, denoted by $L(a)$, which is exogenously given. The latter function, together with the Pareto distribution of firm size, determines the productivity of job positions. Conveniently, we assume that:

$$L(a) = \left(\frac{\pi - m}{\pi}\right)\left(\frac{a}{a_L}\right)^m, \tag{5}$$

where $0 < m < \pi$ (i.e., m is a positive constant with a value lower than the tail index of the Pareto distribution of firm size).

It then follows that the total measure of jobs is one and that the productivity of a firm's position is also Pareto with a tail index of $\pi - m$, i.e., the productivity of a job position with ranking j can be expressed as $a\,[j] = a_L\,(1-j)^{-\gamma}$, where $\gamma \equiv \frac{1}{\pi - m} > 0$. A higher m implies a greater dispersion in the top ranks, in the sense that more productive job positions become more probable at the top. Implicitly, we also assume that there is no shortage in the supply of skilled workers for any available position at a firm, so that the latter can always be filled.

Production function. We assume complementarity between a firm's asset size (a) and its workers' skills (s). Moreover, we assume that a firm's production function is additively separable across its workers' types as well as multiplicatively separable within their types. Consequently, the firm's productivity factor applies multiplicatively to the the aggregate skill of the workers it hires (to fill its positions $L\,(a)$). That is, if $\theta > 0$ is the complementarity parameter, the production function of firm a that hires $l(s)$ number of workers of type s is:

$$Q\,(a,s) = a^{\theta}\left(\int s l\,(s)\,ds\right). \tag{6}$$

Cost of supplying amenities. The second key departure from a standard assignment model is that firms differ in their cost of providing amenities. We assume that the total cost that firm a incurs for providing amenity n at quantity x_n, denoted by $C_n\,(x_n, a)$, is proportional to the quantity x_n. What matters for firm a's profit maximization problem is amenity n's cost *per* position, which is specified as:

$$\frac{C_n\,(x_n, a)}{L\,(a)} = x_n\left(c_n a^{-\kappa_n}\right), \tag{7}$$

where $c_n a^{-\kappa_n}$ represents firm a's marginal cost of providing one additional unit of amenity n per worker.

The parameter κ_n is the central object of interest in our model. When $\kappa_n = 0$, the marginal cost of amenity n is constant across firms — size confers no advantage. When $\kappa_n > 0$, larger firms face a lower marginal cost per position, reflecting economies of scale in amenity provision: this is the competitive advantage channel. When $\kappa_n < 0$, smaller firms have the advantage. As we show below, it is the preference-weighted sum $\sum_{n=1}^{N} \kappa_n \alpha_n$ that determines whether the slope of worker utility exceeds that of wages with firm size — the precise empirical pattern documented in Section 3. This makes κ_n indirectly identifiable from the data: the gap between the elasticity of worker satisfaction and the elasticity of wages with respect to firm size is our empirical signal of competitive advantage.

Firms' profit maximization. The separability of workers' skills in a firm's production function (leading to Eq. (6) above) implies that, in a competitive equilibrium, a firm with productivity or size a (also referred to as firm a) fills its positions by hiring only workers of a particular type s. This implication simplifies to a great extent the competitive equilibrium's characterization. In particular, firm a's profit maximization problem can be written as:

$$V(a) = L(a) \max_{s, \boldsymbol{x}} \left\{ a^{\theta} s - \sum_{n=0}^{N} x_n c_n a^{-\kappa_n} \mid u(\boldsymbol{x}) \geq U(s) \right\}. \tag{8}$$

$U(s)$ indicates the equilibrium utility of workers of type s. The problem is thus equivalent to the problem of maximizing firm a's profit per position (i.e., its average profit), denoted by $\overline{V}(a) \equiv \frac{V(a)}{L(a)}$, by hiring workers of type s and providing them with the amenities they require.

Isomorphism to one-to-one matching. While each firm has multiple positions to fill, the additive separability in Eq. (6) implies that each position contributes independently to the firm's total output— each position's problem is identical and can be solved separately. This means the many-positions-to-one-firm matching problem becomes isomorphic to a one-to-one matching framework, substantially simplifying the equilibrium characterization. As such, the relevant distribution will be that of job positions.[2]

Definition of equilibrium. Given any lower bound of firms' productivity (a_L) and any lower or upper bound of workers' skills (s_L or s_H), a competitive labor market equilibrium

[2]Recall that the distribution of the productivity of a firm's job position is Pareto with a tail index of $1/\gamma = \pi - m$.

consists of (i) the equilibrium utility $U(s)$ of workers with skills s, and (ii) an assignment function $\sigma(a)$ specifying the optimal type of workers (i.e., the index of their skills) hired by a firm with productivity or size a, such that (iii) every firm maximizes its profit per position (so that Eq. (8)'s optimality conditions hold), and (iv) the market-clearing condition is satisfied for the labor market.

4.2. Sorting

The equilibrium determines jointly a firm's wage and amenities, as well as the type of its workers. We now analyze the sorting pattern, and the corresponding compensation packages. According to Eq. (8), firm a chooses to provide the compensation bundle $\boldsymbol{x} = (x_0, x_1, ..x_m)$ such that it minimizes its expenditure per position subject to the participation constraint of worker s. That is, the overall packaging offered must give worker s his equilibrium utility $U(s)$. The solution of the compensation bundle can be interpreted as a Hicksian demand function given the amenities' marginal costs per position that firm a incurs and the equilibrium utility $U(s)$ of worker s. Specifically, the quantity of amenity n that firm a offers to worker s is given by:

$$x_n(a,s) = \frac{\alpha_n}{c_n a^{-\kappa_n}} \left(\psi \, a^{-\sum_{n=1}^{N}(\kappa_n \alpha_n)} \, U(s) \right). \tag{9}$$

where, for convenience, we define the constant $\psi \equiv \prod_{n=0}^{N} \left(\frac{c_n}{\alpha_n}\right)^{\alpha_n}$.[3]

Equation (9) reveals the mechanism through which competitive advantage shapes compensation. For a given worker utility level $U(s)$, a firm with a larger κ_n faces a lower cost of providing amenity n and therefore supplies more of it. This is not merely a repackaging of the same total compensation — it is a genuine efficiency gain that allows larger firms to deliver higher utility per dollar of expenditure.

Given the equilibrium utility $U(s)$ of worker s, firm a chooses the worker s^* that solves:

$$\overline{V}(a) = \max_{s} \left\{ a^{\theta} s - \psi \, a^{-\sum_{n=1}^{N}(\kappa_n \alpha_n)} \, U(s) \right\}, \tag{10}$$

where the second term, $\psi \, a^{-\sum_{n=1}^{N}(\kappa_n \alpha_n)} \, U(s)$, represents the average expenditure per position

[3] Recall that wage can be treated as a special case of amenity where the monetary costs to all firms are the same (i.e., $\kappa_0 = 0$ and $c_0 = 1$), and thus $x_0(a,s) = \alpha_0 \left(\psi \, a^{-\sum_{n=1}^{N}(\kappa_n \alpha_n)} \, U(s) \right)$.

that firm a incurs to hire worker s. Since ψ is constant across firms, it can be read as a uniform cost factor, whereas $a^{-\sum_{n=1}^{N}(\kappa_n\alpha_n)}$ is the firm-specific cost factor conditional on the worker's equilibrium utility.

Eq. (10) highlights that firm size has two effects: The first one is the standard complementarity in the production function, captured by the parameter θ: larger firms are more productive and therefore willing to pay more for skilled workers. The second is the cost advantage in supplying amenities, summarized by $\sum_{n=1}^{N}(\kappa_n\alpha_n)$: when this term is positive, larger firms can deliver a given utility level at lower cost, allowing them to bid more aggressively for talent. Both forces push in the same direction when $\sum_n \kappa_n\alpha_n \geq 0$, but they can work against each other when smaller firms have the amenity cost advantage. The sorting outcome is determined by which force dominates.

Lemma 1. *The larger the size of a firm, the higher the skill of the workers it hires, i.e., there is positive assortative matching (PAM), if and only if* $\theta + \sum_{n=1}^{N}(\kappa_n\alpha_n) > 0$.

Lemma 1 establishes the condition for PAM. When $\sum_{n=1}^{N}(\kappa_n\alpha_n) \geq 0$, PAM is always satisfied — larger firms both benefit more from hiring skilled workers and face lower costs in providing amenities, so they unambiguously dominate in the competition for talent. When $\sum_{n=1}^{N}(\kappa_n\alpha_n) < 0$, these two forces counteract each other: larger firms benefit more from skilled workers through the production function, but face higher amenity costs. PAM continues to hold only if the production complementarity dominates, i.e., $\theta + \sum_{n=1}^{N}(\kappa_n\alpha_n) > 0$. Our empirical estimates in Section 5 confirm that this condition is satisfied, consistent with the positive assortative matching between firm size and worker quality observed in the data.

4.3. The Role of Competitive Advantage

We now show how competitive advantage generates the slope differential between worker utility and wages that we documented in Section 3. This is the key theoretical result of the paper: it establishes that the differential slopes of satisfaction and wages with firm size are not just consistent with competitive advantage, but are a precise empirical signal of it — one we can use to quantify the relative contributions of selection and competitive advantage in Section 5.

Complementarity between competitive advantage and workers' preferences. The utility schedule $U[i]$ depends on market competition. The first-order condition of (10) implies

$$\psi\, U'[i] = s'[i]\big(a[i]\big)^{\theta+\sum_{n=1}^{N}\kappa_n\alpha_n}. \tag{11}$$

Thus, the marginal utility of workers depends both on their marginal contribution to output and on firms' cost advantage in providing amenities. When firms have no competitive advantage ($\kappa_n = 0$ for all n), the slope reduces to that of a standard assignment model:[4]

$$U'[i] = s'[i]\big(a[i]\big)^{\theta}.$$

In this case, utility and wages scale identically with firm size — there is no slope differential. The slope differential emerges only when $\sum_{n=1}^{N}\kappa_n\alpha_n > 0$, i.e., when larger firms have a cost advantage in providing amenities that workers value. This is the model's central identification insight: the gap between the elasticity of overall ratings and the elasticity of wages with respect to firm size directly measures $\sum_{n=1}^{N}\kappa_n\alpha_n$.

The effective role of firm size is summarized by $\theta + \sum_{n=1}^{N}\kappa_n\alpha_n$, which captures both production complementarity and the interaction between competitive advantage and workers' preferences. A firm's advantage in providing amenity n (κ_n) is more valuable when workers place greater weight on that amenity (higher α_n). The term $\sum_{n=1}^{N}\kappa_n\alpha_n$ therefore summarizes the value of firms' competitive advantage.

On workers' utilities. Equation (11) implies that a higher $\theta + \sum_{n=1}^{N}\kappa_n\alpha_n$ leads to a steeper utility schedule across rankings. Intuitively, firms become more heterogeneous, so the gains from matching with larger firms increase —both because they are more productive and because they provide amenities more efficiently. This heterogeneity is what drives the utility premium at superstar firms beyond what their wage premium alone would predict.

To obtain a closed-form expression for equilibrium utility, we impose the following assumption, which facilitates the calibration in Section 5. The equilibrium utility $U[i]$ is pinned down up to the constant $U[0]$. All results continue to hold if this assumption is relaxed and attention is restricted to the upper tail, as in Gabaix and Landier (2008).

Assumption 1.

$$U[0] = \frac{B}{\gamma\Big(\theta + \sum_{n=1}^{N}\kappa_n\alpha_n\Big) + \beta}\left(\frac{1}{\psi}\right) a_L^{\theta+\sum_{n=1}^{N}\kappa_n\alpha_n}.$$

[4]This implies that even though workers value amenities, the slope of equilibrium utility coincides with that in Gabaix and Landier (2008) when firms lack a competitive advantage.

Proposition 1. *Under Assumption 1 and PAM, equilibrium utility satisfies*

$$U[i] = \left(\frac{B}{\gamma\left(\theta + \sum_{n=1}^{N} \kappa_n \alpha_n\right) + \beta} \left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}} \right) \left(\frac{1}{\psi}\right) \left(a[i]\right)^{\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma}}. \tag{12}$$

Equilibrium utilities therefore follow a Pareto distribution with tail index $\frac{1}{\left(\theta + \sum_{n=1}^{N} \kappa_n \alpha_n\right) + \frac{\beta}{\gamma}}$.

Proposition 1 delivers a log-linear relationship between worker utility and firm size:

$$\ln U[i] = u_0 + \left(\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma} \right) \ln a[i], \tag{13}$$

where u_0 is a constant. The slope of this relationship, $\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma}$, has three components: the production complementarity θ, the competitive advantage term $\sum_{n=1}^{N} \kappa_n \alpha_n$, and the worker talent heterogeneity term $\frac{\beta}{\gamma}$. By contrast, the slope of wages with firm size (derived below) is $\theta + \frac{\beta}{\gamma}$ — the same expression without the competitive advantage term. The slope differential between utility and wages is therefore exactly $\sum_{n=1}^{N} \kappa_n \alpha_n$, which is zero when firms have no cost advantage in amenity provision and positive when they do. This gives us a clean empirical strategy: estimate the slope of wages and the slope of overall ratings with respect to firm size, and attribute the difference to competitive advantage.

Compensation bundles. The characterization of workers' equilibrium compensation bundles follows directly from the equilibrium utility. Substituting (12) into (9) yields the equilibrium quantity of amenity n:

$$x_n[i] = \left(\frac{B}{\gamma\left(\theta + \sum_{n=1}^{N} \kappa_n \alpha_n\right) + \frac{\beta}{\gamma}} \left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}} \right) \frac{\alpha_n}{c_n} \left(a[i]\right)^{\theta + \frac{\beta}{\gamma} + \kappa_n}. \tag{14}$$

The quantile function of wages is therefore Pareto with tail index $\frac{1}{\gamma\left(\theta + \frac{\beta}{\gamma}\right)}$, while the quantile function of amenity n is Pareto with tail index $\frac{1}{\gamma\left(\theta + \kappa_n + \frac{\beta}{\gamma}\right)}$.

Corollary 1. *Under Assumption 1 and PAM,*

$$\ln x_n[i] = \eta_n - \frac{\beta}{\gamma} \ln a[i^*] + \left(\theta + \frac{\beta}{\gamma} + \kappa_n \right) \ln a[i], \tag{15}$$

where η_n is a constant. Since wages correspond to the special case $\kappa_n = c_n = 0$, we have

$$\ln w[i] = \eta_0 - \frac{\beta}{\gamma} \ln a[i^*] + \left(\theta + \frac{\beta}{\gamma} \right) \ln a[i]. \tag{16}$$

Corollary 1 makes the identification strategy precise. Equations (15) and (16) show that wages and amenities scale with firm size according to the same base rate $\theta + \frac{\beta}{\gamma}$, but amenity n has an additional term κ_n in its slope. Overall worker utility, which aggregates wages and all amenities, therefore scales at rate $\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma}$ — faster than wages if and only if $\sum_{n=1}^{N} \kappa_n \alpha_n > 0$. Equation (16) implies that $\theta + \frac{\beta}{\gamma} > 0$ is consistent with the positive relationship between wages and firm size observed in the data. Even when $\kappa_n = 0$, compensation rises with firm size due to selection if $\theta + \frac{\beta}{\gamma} > 0$; competitive advantage amplifies this by increasing amenity provision relative to wages.

Together, these equations map directly onto the two regression specifications in Section 3: Eq. (16) corresponds to the wage regression in Panel A of Table 4, while Eq. (13) — via the rating proxy of Sockin (2024) — corresponds to the overall rating regression in Panel C. The difference in estimated slopes between the two panels is our empirical measure of $\sum_{n=1}^{N} \kappa_n \alpha_n$, which we recover structurally in Section 5.

5. Decomposing Selection versus Competitive Advantage

The model delivers a sharp empirical test. From Equations (16) and (13), wages and worker utility both scale log-linearly with firm size, but with different slopes. The slope of wages is $\theta + \frac{\beta}{\gamma}$, capturing production complementarity and worker talent heterogeneity — the pure selection channel. The slope of utility exceeds that of wages by exactly $\sum_{n=1}^{N} \kappa_n \alpha_n$, the preference-weighted competitive advantage term. If firms had no cost advantage in amenity provision ($\kappa_n = 0$ for all n), utility and wages would scale identically with firm size and we would observe no slope differential. The fact that we do — with an elasticity of 0.079 for overall ratings versus 0.045 for wages — is therefore direct evidence of competitive advantage.

This observation suggests a clean identification strategy. We estimate the wage equation (16) to recover the selection parameters θ and β/γ. We then estimate the ratings equation (17) to recover the competitive advantage parameter $\sum_{n=1}^{N} \kappa_n \alpha_n$, using the slope differential between ratings and wages as the identifying variation. Crucially, this approach requires only data on wages and overall firm ratings — it does not require direct measurement of amenities, which are not systematically observed. All calibration results are summarized in Table 5.

5.1. Inferring from the Utility-Wage Slope Differential

Workers' utilities and overall ratings. Since we do not observe workers' utilities directly, we follow Sockin 2024 and model workers' overall firm ratings as a noisy proxy for

Table 5
Summary of the model's calibration

This table summarizes the calibration of the model. Column 1 shows the parameters. Column 2 shows their estimated value. Column 3 shows the data and equation based on which the parameters' estimation is performed. See Table 1 for a detailed description of the sample.

(1) Parameters	(2) Values	(3) Firm-level Data and Equations
θ	0.145	Firm *Wage* and *Assets*
λ	0.710	Firm *TotRating* and *Assets*
$\sum_n \kappa_n \alpha_n$	0.066	Firm *TotRating* and *Assets*
π	1.158	Firm *Assets*
m	0.662	Firm *EMP* and *Assets*
γ	0.496	$\hat{\gamma} = \hat{\pi} - \hat{m}$
β	−0.050	$\hat{\beta} = \hat{\gamma} \cdot \widehat{\beta/\gamma}$
α_0	0.818	*IB* per worker to model-adjusted *Assets* & model-adjusted *Wage* to assets ratio
	0.790	*NI* per worker to model-adjusted *Assets* & model-adjusted *Wage* to assets ratio

utility:

$$TotRating[i] = \exp(\epsilon)\, U[i]^{\lambda},$$

where $\lambda > 0$ and ϵ is an error term. The parameter λ captures the sensitivity of observed ratings to underlying utility — it need not equal one, since ratings are measured on a bounded five-star scale rather than in utility units. Substituting Proposition 1 into this expression yields a testable equation linking observed ratings to firm size.

Corollary 2. *Consider a position in a reference firm (e.g., the median firm) with rank i^*. Then*

$$\ln TotRating[i] = \zeta - \lambda \frac{\beta}{\gamma} \ln a[i^*] + \lambda \left(\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma} \right) \ln a[i] + \epsilon, \tag{17}$$

where ζ is a constant.

Comparing Equations (16) and (17) clarifies the identification. The coefficient on firm assets in the wage regression identifies $\theta + \frac{\beta}{\gamma}$ — the selection component. The coefficient on firm assets in the ratings regression identifies $\lambda(\theta + \sum_n \kappa_n \alpha_n + \frac{\beta}{\gamma})$ — the total utility scaling, which includes both selection and competitive advantage. The coefficient on industry median assets in the ratings regression identifies $\lambda \frac{\beta}{\gamma}$, which, combined with the wage regression estimate of $\frac{\beta}{\gamma}$, recovers λ. With λ in hand, $\sum_n \kappa_n \alpha_n$ is identified as the residual — the portion of utility scaling with firm size that cannot be attributed to selection alone.

5.2. Decomposition Results

From Column 3 of Panel A of Table 4, the estimated coefficient on $\ln(Assets_{i,t})$ is 0.045 and the coefficient on $\ln(IndustryMedianAssets_{i,t})$ is 0.100. These estimates imply $\beta/\gamma = -0.100$ and $\theta = 0.145$. The negative value of β/γ is consistent with the $\beta < 0$ case in Equation (3), indicating that the distribution of worker talent has an upper bound — workers at the top of the skill distribution are relatively similar to each other, which will matter for the welfare analysis in Section 7. The estimate $\theta = 0.145$ implies decreasing returns to scale in production, consistent with the standard assignment model literature.

Turning to Equation (17), Column 3 of Panel C shows that the coefficient on firm assets is 0.079 and that on median industry assets is 0.071. Dividing the latter by $|\beta/\gamma|$ yields $\lambda \approx 0.710$, confirming that observed ratings are a somewhat compressed measure of underlying utility. Using this estimate, we infer $\sum_{n=1}^{N} \kappa_n \alpha_n \approx 0.066$.

The key decomposition follows from Equation (13). The total scaling of worker utility with firm size is governed by $\theta + \sum_n \kappa_n \alpha_n + \frac{\beta}{\gamma} = 0.045 + 0.066 = 0.111$. Of this, competitive advantage contributes:

$$\frac{\sum_{n=1}^{N} \kappa_n \alpha_n}{\theta + \sum_{n=1}^{N} \kappa_n \alpha_n + \frac{\beta}{\gamma}} = \frac{0.066}{0.045 + 0.066} \approx 60\%,$$

while selection accounts for the remaining 40%.

This is the central quantitative result of the paper. Even if superstar firms had no cost advantage in amenity provision — if $\kappa_n = 0$ for all n — they would still be better places to work than smaller firms, because they attract more talented workers who demand higher compensation in all its forms. But this selection effect accounts for only 40% of the observed utility premium. The majority of the premium — 60% — reflects a genuine structural efficiency: superstar firms can produce workplace utility more cheaply per worker than their smaller competitors, and they optimally exploit this advantage by substituting amenities for cash wages. This substitution raises profits while maintaining or increasing worker utility, a connection we formalize in the next section.

6. Implications for Corporate Profits

The decomposition in Section 5 established that competitive advantage in amenity provision accounts for roughly 60% of the utility premium at superstar firms. We now examine the flip side of this result: if superstar firms can deliver a given level of worker utility more cheaply than their competitors, this efficiency must show up in profits. We show that workers' amenity preferences affect profit sharing if and only if firms possess a competitive advantage ($\kappa_n > 0$) — and that the quantitative impact is economically significant.

Under PAM, a worker with skill ranking i is matched to a firm position with productivity ranking i. Firm i's profit per position can be written as

$$\overline{V}[i] = s[i]\big(a[i]\big)^{\theta} - \psi\left(a[i]\right)^{-\sum_{n=1}^{N}\kappa_n\alpha_n}U[i], \tag{18}$$

where $U[i] \equiv U(s[i])$ denotes the equilibrium utility promised to worker i. The first term is revenue per position, which depends only on firm size and worker skill. The second term is the firm's expenditure per position, which depends on both the worker's equilibrium utility and the firm's cost factor $\psi\left(a[i]\right)^{-\sum_n \kappa_n\alpha_n}$. When $\sum_n \kappa_n\alpha_n > 0$, this cost factor is decreasing in firm size — larger firms can deliver a given utility level at lower cost — and profits are correspondingly higher.

Firm heterogeneity and corporate profits. To understand how amenity preferences affect the distribution of profits, note that equilibrium expenditure is given by

$$e[i] = \left(\frac{B}{\gamma\Big(\theta + \sum_{n=1}^{N}\kappa_n\alpha_n\Big) + \beta}\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)\big(a[i]\big)^{\theta+\frac{\beta}{\gamma}}. \tag{19}$$

A key feature of Equation (19) is that competitive advantage — $\sum_n \kappa_n\alpha_n$ — enters the prefactor but not the exponent on $a[i]$. This means that competitive advantage raises firm heterogeneity (through the denominator of the prefactor), compressing expenditure uniformly across firms rather than disproportionately at the top. As workers place greater weight on amenities that larger firms supply more cheaply, those firms become harder to replace as employers — they are the only ones who can efficiently provide what workers want. This strengthens their bargaining position in the labor market, allowing them to capture a greater share of the matching surplus.

Proposition 2. *A higher α_n reduces expenditure and increases profits for all firms under PAM if and only if $\kappa_n > 0$. In particular,*

$$\frac{\partial \ln e[i]}{\partial \ln \alpha_n} = -\frac{\kappa_n\alpha_n}{\theta + \sum_{n=1}^{N}\kappa_n\alpha_n + \frac{\beta}{\gamma}}. \tag{20}$$

The if-and-only-if condition in Proposition 2 is important: amenity preferences only translate into higher profits when firms have a genuine cost advantage in providing those amenities ($\kappa_n > 0$). If $\kappa_n = 0$ — amenities are equally costly for all firms — then rising amenity demand simply reshuffles the compensation bundle without affecting total expenditure or profits. It is the combination of worker demand and firm-side cost efficiency that generates the profit effect. When $\kappa_n < 0$, the mechanism reverses: higher amenity preferences favor smaller firms, increasing their bargaining power and raising expenditure at larger firms.

Economic significance. Using our calibrated parameters from Table 5, Equation (20) implies that a 1% increase in amenity preferences reduces firm expenditure by approximately 0.6%, generating a commensurate increase in profits. To put this in perspective, our estimate of $\sum_n \kappa_n \alpha_n \approx 0.066$ suggests that the demand for non-pecuniary amenities is already generating meaningful profit gains for superstar firms — and as amenity demand continues to grow, the profit advantage of firms with scale-based cost efficiency will widen further.

Note that this elasticity does not depend on the absolute weight workers place on amenities $(1 - \alpha_0)$ in any first-order way. The profit impact is driven by the interaction between workers' marginal preferences and firms' cost advantage: even if the total weight on amenities is modest, large differences in κ_n across firm sizes translate into large marginal profit effects. We discuss how to pin down α_0 and other auxiliary parameters in the Appendix B.

These findings establish workplace culture as a profit-generating intangible — not merely a cost of attracting talent, but a competitive lever that superstar firms can exploit more effectively than their rivals. As workers increasingly demand non-pecuniary benefits, firms with scale advantages will optimally expand amenity provision at the expense of cash wages, reducing the measured labor share while raising profits. It is to this labor share implication that we now turn.

7. Implications for the Falling Labor Share

7.1. Role of Increasing Demand for Workplace Amenities

A secular decline in the labor share, particularly among high-productivity "superstar firms," is one of the most debated facts in modern macroeconomics (Autor et al. 2020). The prevailing narrative attributes this decline to rising market power: as superstar firms capture greater market share, they use their bargaining advantage to suppress wages, returning a shrinking fraction of value added to workers. Under this view, a higher γ — greater market concentration — reduces both the labor share and worker welfare simultaneously. The two move together by construction, and a declining labor share is an unambiguous signal of deteriorating worker outcomes.

We offer a complementary mechanism in which this link breaks down. When workers increasingly demand non-pecuniary amenities (Mas and Pallais 2017, He et al. 2021), a higher α_n reduces the labor share through the competitive advantage channel — but can simultaneously raise worker welfare. The labor share and worker utility need not move in the same direction, and whether they do depends on which channel is driving the decline. We do not attempt to separately identify the market power and amenity demand channels empirically; our point is that rising amenity demand alone is capable of producing a declining labor share, and that when it does, the standard interpretation of that decline as evidence of worker harm need not apply.

There are two distinct forces through which our mechanism reduces the labor share. The first is an accounting channel: the *pecuniary* labor share recorded in national accounts is partially offset by a rise in an *intangible labor share* composed of non-pecuniary workplace amenities. Conventional labor-share measures therefore systematically understate the true compensation of workers at highly productive firms. To formalize this distinction, we separate the *pecuniary labor share*, which is recorded in national accounts, and the *total labor share*, which accounts for firms' expenditures on workplace culture, amenities, and other non-pecuniary benefits.

Formally, for a firm at rank i,

$$LS_w[i] \equiv \frac{x_0[i] \cdot L[i]}{Q[i]}, \qquad LS_e[i] \equiv \frac{e[i] \cdot L[i]}{Q[i]}. \tag{21}$$

Since only a fraction α_0 of total compensation is paid in cash wages, the two measures are related by $LS_w[i] = \alpha_0 \, LS_e[i]$. This immediately implies that part of the measured decline in the labor share may reflect a *compositional shift* toward non-pecuniary compensation rather than a true reduction in workers' total compensation. As workers place greater weight on amenities and workplace culture, firms optimally substitute away from cash wages — lowering the measured labor share even if total compensation remains high. A rise in α_n necessarily implies a fall in $\alpha_0 = 1 - \sum_n \alpha_n$, mechanically lowering the measured wage share and further amplifying the observed decline.

The second force operates through the total labor share itself. Even accounting for all expenditure, the total labor share can decline when workers' non-pecuniary preferences increase — but only when firms possess a competitive advantage in amenity provision.

Corollary 3. *The total labor share decreases in firm rank i (when $\beta < 0$) and in the Pareto* *index π. It is decreasing in workers' non-pecuniary preferences α_n if and only if $\kappa_n > 0$:*

$$LS_e[i] = \frac{-\beta}{\gamma(\theta + \sum_n \kappa_n \alpha_n) + \beta} \cdot \frac{t(i)}{s_H - t(i)}, \qquad t(i) \equiv \left(\frac{1}{\tilde{B}}\right)^{-\beta} (1-i)^{-\beta}. \tag{22}$$

When $\beta < 0$, $t(i)$ is decreasing in rank i, implying lower labor shares among larger, more productive firms. This is consistent with the existing literature: superstar firms exhibit lower labor shares, and labor shares decline with market concentration (Autor et al. 2020). When firms possess a competitive advantage in supplying amenities ($\kappa_n > 0$), an increase in workers' preference for these amenities effectively raises the value of the firm's technology. This allows firms to extract greater surplus and hire workers at lower total expenditure cost. As a result, the total labor share declines — even though workers may be better off in utility terms. The if-and-only-if condition in Corollary 3 mirrors that in Proposition 2: it is precisely when competitive advantage is present that rising amenity demand tightens the labor market in favor of superstar firms, compressing expenditure and reducing the labor share.

7.2. Divergence of Labor Share and Worker Welfare

The most important distinction between our mechanism and the market power narrative lies in the welfare implications. Under the market power channel, a declining labor share and deteriorating worker welfare move together — firms extract more rent, workers receive less, and the two are inseparable. Our competitive advantage channel breaks this link: an increase in α_n can simultaneously produce a decline in the labor share and a rise in worker utility. We now formalize this divergence.

7.2.1. Cross-Sectional Implications: The Utility–Expenditure Gap

This decoupling is already visible at the cross-sectional level. Worker utility satisfies

$$\ln U[i] = -\ln \psi + \ln e[i] + \sum_{n=1}^{N} \kappa_n \alpha_n \ln a[i]. \tag{23}$$

Define the utility–expenditure gap as $\ln U[i] - \ln e[i]$. Whenever $\sum_{n=1}^{N} \kappa_n \alpha_n > 0$, this gap is increasing in firm rank i, since $a[i]$ is increasing in firm size and productivity. The implication is that the total labor share systematically understates workers' utility more for larger firms — precisely the firms where the market power narrative predicts the greatest welfare losses. While superstar firms exhibit lower expenditure shares on labor, workers' effective standard of living rises with firm size due to the scaling of amenities with productivity. How worker utility responds to a marginal increase in amenity preferences is given by:

$$\frac{\partial \ln U[i]}{\partial \alpha_n} = \kappa_n \left(\ln a[i] - \frac{\gamma}{\gamma \left(\theta + \sum_{m=1}^{N} \kappa_m \alpha_m \right) + \beta} \right) + \ln\left(\frac{\alpha_n}{c_n}\right) + 1.$$

When $\kappa_n > 0$, this expression is increasing in firm size $a[i]$, confirming that the utility gains from rising amenity demand are concentrated at the top of the firm-size distribution — precisely where the labor share is lowest and the amenity scaling is strongest.

7.2.2. Effect of Non-pecuniary Preferences on Average Worker Welfare

Whether the divergence between labor shares and welfare holds in the aggregate depends on the balance of three competing forces. Define $\Theta \equiv \theta + \sum_{n=1}^{N} \kappa_n \alpha_n$. Average utility is given by

$$\bar{U} = \frac{B a_L^{\Theta}}{\psi(\gamma\Theta + \beta)\big(1 - (\gamma\Theta + \beta)\big)}, \tag{24}$$

provided $\gamma\Theta + \beta < 1$. An increase in α_n affects average utility through three channels.

Fat-tail effect. A higher Θ increases the thickness of the upper tail of the firm-size distribution, captured by the term $1-(\gamma\Theta+\beta)$ in the denominator. As the tail fattens, top-ranked workers benefit disproportionately, raising mean utility.

Productivity effect. A higher Θ increases the productivity scaling term a_L^{Θ} when $a_L > 1$. As firms' competitive advantage becomes more valuable, larger firm size creates higher matching surplus, increasing the total pie available to both workers and firms.

Profit-sharing effect. A higher Θ increases $\gamma\Theta + \beta$, shifting surplus toward firms and reducing workers' share of output, which lowers average utility. This is the channel that the market power narrative emphasizes exclusively — and it always operates in the direction of lower welfare.

The first two channels raise worker welfare while the third reduces it. The net effect yields

$$\frac{\partial \ln \bar{U}}{\partial \alpha_n} = 1 - \ln\left(\frac{c_n}{\alpha_n}\right) + \kappa_n \left[\ln a_L - \gamma \frac{1 - 2(\gamma\Theta + \beta)}{(\gamma\Theta + \beta)\big(1 - (\gamma\Theta + \beta)\big)} \right]. \tag{25}$$

The last term is positive if $\kappa_n > 0$ and firm scale is large enough — that is, if the productivity and fat-tail effects dominate the profit-sharing effect. This condition is not guaranteed in general; it requires that the economy be sufficiently large that the gains from higher matching surplus outweigh the loss of workers' bargaining power.

Proposition 3. *For any $\kappa_n > 0$, an increase in α_n always reduces average expenditure and thus the labor share, but increases average utility when firm scale a_L is sufficiently large.*

With our calibrated parameters, the condition becomes

$$\ln a_L > \gamma \frac{1 - 2(\gamma\Theta + \beta)}{(\gamma\Theta + \beta)\big(1 - (\gamma\Theta + \beta)\big)} \approx 8.6,$$

that is, $a_L > 5000$. The threshold condition $a_L > 5000$ is satisfied in our empirical setting, which focuses on large S&P 500 firms. Intuitively, when the economy is populated by sufficiently large firms, the productivity and fat-tail effects — which raise welfare by increasing matching surplus and fattening the upper tail of the firm-size distribution — dominate the profit-sharing effect that shifts surplus toward firms. As a result, rising demand for workplace amenities generates a declining labor share and rising worker utility simultaneously. This is the sense in which our mechanism offers a fundamentally different interpretation of the superstar labor share decline: not as evidence of worker exploitation, but as the equilibrium consequence of workers and firms efficiently reallocating compensation toward the non-pecuniary benefits that modern workers increasingly value.

8. Conclusion

Why are superstar firms the best places to work? By documenting that worker satisfaction scales nearly twice as fast as wages with firm size, we identify a "slope differential" that cannot be explained by talent sorting alone. Our assignment model reveals that superstar firms possess a unique competitive advantage: the ability to leverage scale to provide non-pecuniary amenities more efficiently than cash wages. We find that this efficiency channel explains roughly three-fifths of the utility premium associated with firm size. Consequently, superstar firms do not simply pay less relative to their productivity; they pay differently, substituting workplace utility for cash in a way that raises both firm profits and worker satisfaction.

These findings offer a new perspective on the secular decline of the labor share. The prevailing narrative — that declining labor shares reflect rising market power and imply lower worker welfare — rests on the assumption that labor shares and worker utility move together. Our results show that this assumption breaks down when firms compete on non-pecuniary amenities. As workers increasingly demand workplace culture, flexibility, and other non-pecuniary benefits, firms with a scale-based cost advantage optimally substitute away from cash wages, mechanically reducing the measured labor share. Whether workers are better or worse off depends on the balance of three forces — a profit-sharing effect that favors firms, and productivity and fat-tail effects that favor workers — and on a condition on firm scale that must be satisfied for the worker-welfare result to hold. Given our estimated parameters, that condition is met, and the average worker is better off. The superstar economy need not be one of rent extraction; it may instead be one in which the most productive firms are increasingly specialized in producing the intangible cultural assets that modern workers value most.

References

Abowd, John M, Francis Kramarz, and David N Margolis (1999). “High wage workers and high wage firms”. In: *Econometrica* 67.2, pp. 251–333.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen (2020). “The Fall of the Labor Share and the Rise of Superstar Firms”. In: *The Quarterly Journal of Economics* 135.2, pp. 645–709.

Brown, Charles and James Medoff (1989). “The employer size-wage effect”. In: *Journal of Political Economy* 97.5, pp. 1027–1059.

Chang, Briana and Harrison Hong (2019). “Selection versus talent effects on firm value”. In: *Journal of Financial Economics* 133.3, pp. 751–763.

— (2026). *Impact Trickles Down: A General Equilibrium Theory of Stakeholder Exit and Engagement*. Tech. rep. National Bureau of Economic Research.

Chiappori, Pierre-André and Bernard Salanié (2016). “The econometrics of matching models”. In: *Journal of Economic Literature* 54.3, pp. 832–61.

Colonnelli, Emanuele, Timothy McQuade, Gabriel Ramos, Thomas Rauter, and Olivia Xiong (2024). “Polarizing corporations: Does talent flow to “good” firms?” In: *Working Paper*.

Dehaan, Ed, Nan Li, and Frank S Zhou (2023). “Financial reporting and employee job search”. In: *Journal of Accounting Research* 61.2, pp. 571–617.

Edmans, Alex (2011). “Does the stock market fully value intangibles? Employee satisfaction and equity prices”. In: *Journal of Financial Economics* 101.3, pp. 621–640.

Edmans, Alex, Darcy Pu, Chendi Zhang, and Lucius Li (2024). “Employee satisfaction, labor market flexibility, and stock returns around the world”. In: *Management Science* 70.7, pp. 4357–4380.

Eisfeldt, Andrea L, Antonio Falato, and Mindy Z Xiaolan (2023). “Human Capitalists”. In: *NBER Macroeconomics Annual* 37.1, pp. 1–61.

Gabaix, Xavier and Rustam Ibragimov (2011). “Rank- 1/2: a simple way to improve the OLS estimation of tail exponents”. In: *Journal of Business & Economic Statistics* 29.1, pp. 24–39.

Gabaix, Xavier and Augustin Landier (2008). “Why has CEO pay increased so much?” In: *The Quarterly Journal of Economics* 123.1, pp. 49–100.

Green, T Clifton, Ruoyan Huang, Quan Wen, and Dexin Zhou (2019). “Crowdsourced employer reviews and stock returns”. In: *Journal of Financial Economics* 134.1, pp. 236–251.

Grossman, Gene M and Ezra Oberfield (2022). “The elusive explanation for the declining labor share”. In: *Annual Review of Economics* 14.1, pp. 93–124.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales (2015). “The Value of Corporate Culture”. In: *Journal of Financial Economics* 117.1, pp. 60–76.

He, Haoran, David Neumark, and Qian Weng (2021). “Do workers value flexible jobs? A field experiment”. In: *Journal of Labor Economics* 39.3, pp. 709–738.

Huang, Kelly, Meng Li, and Stanimir Markov (2020). “What do employees know? Evidence from a social media platform”. In: *The Accounting Review* 95.2, pp. 199–226.

Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed (1998). “Hedonic wages and labor market search”. In: *Journal of Labor Economics* 16.4, pp. 815–847.

Hwang, Hae-shin, W. Robert Reed, and Carlton Hubbard (1992). “Compensating wage differentials and unobserved productivity”. In: *Journal of Political Economy* 100.4, pp. 835–858.

Karabarbounis, Marios and Santiago Pinto (2018). “What can we learn from online wage postings? Evidence from Glassdoor”. In: *Economic Quarterly* 104.4Q, pp. 173–189.

Lavetti, Kurt (2023). “Compensating wage differentials in labor markets: Empirical challenges and applications”. In: *Journal of Economic Perspectives* 37.3, pp. 189–212.

Li, Kai, Feng Mai, Rui Shen, and Xinyan Yan (2021). “Measuring corporate culture using machine learning”. In: *The Review of Financial Studies* 34.7, pp. 3265–3315.

Liu, Tim, Christos A Makridis, Paige Ouimet, and Elena Simintzi (2023). “The distribution of nonwage benefits: maternity benefits and gender diversity”. In: *The Review of Financial Studies* 36.1, pp. 194–234.

Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger (2023). “The value of working conditions in the United States and implications for the structure of wages”. In: *American Economic Review* 113.7, pp. 2007–2047.

Mas, Alexandre and Amanda Pallais (2017). “Valuing alternative work arrangements”. In: *American Economic Review* 107.12, pp. 3722–3759.

Mueller, Holger M, Paige P Ouimet, and Elena Simintzi (2017). “Within-firm pay inequality”. In: *The Review of Financial Studies* 30.10, pp. 3605–3635.

Ouimet, Paige and Geoffrey Tate (2023). *Firms with benefits? Nonwage compensation and implications for firms and labor markets*. Tech. rep. National Bureau of Economic Research.

Pan, Yihui (2017). “The determinants and impact of executive-firm matches”. In: *Management Science* 63.1, pp. 185–200.

Rosen, Sherwin (1981). “The economics of superstars”. In: *American Economic Review* 71.5, pp. 845–858.

Rosen, Sherwin (1986). *The theory of equalizing differences*. In: Ashenfelter, Orley C., Layard, Richard (Eds.), Handbook of Labor Economics, Volume 1 (Chapter 12). North Holland, Elsevier, 641–692.

Sattinger, Michael (1993). "Assignment models of the distribution of earnings". In: *Journal of Economic Literature* 31.2, pp. 831–880.

Sockin, Jason (2024). "Show me the amenity: Are higher-paying firms better all around?" In: *Working Paper*.

Sockin, Jason and Aaron Sojourner (2023). "What's the inside scoop? challenges in the supply and demand for information on employers". In: *Journal of Labor Economics* 41.4, pp. 1041–1079.

Sørensen, Morten (2007). "How smart is smart money? A two-sided matching model of venture capital". In: *The Journal of Finance* 62.6, pp. 2725–2762.

Sorkin, Isaac (2018). "Ranking firms using revealed preference". In: *The Quarterly Journal of Economics* 133.3, pp. 1331–1393.

Tervio, Marko (2008). "The difference that CEOs make: An assignment model approach". In: *American Economic Review* 98.3, pp. 642–668.

Whited, Toni (2019). "*JFE* special issue on labor and finance". In: *Journal of Financial Economics* 133.3, pp. 539–540.

Appendix

A. Proofs

A.1. Proof of Lemma 1

Proof. Given workers' equilibrium utility $U(s)$, Eq. (10) can be rewritten as:

$$\overline{V}(a) = a^{-\sum_{n=1}^{N}(\kappa_n \alpha_n)} \max_s \left\{ a^{\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n)} s - \psi\, U(s) \right\}. \tag{A.1}$$

The first-order condition with respect to s is:

$$a^{\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n)} - \psi\, U'(s) = 0. \tag{A.2}$$

Defining the function $F(a, s) = a^{\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n)} - \psi\, U'(s)$, we see that $\frac{\partial F}{\partial s}(a, s) = -\psi\, U''(s) < 0$ (since $U(s)$ is convex) and $\frac{\partial F}{\partial a}(a, s) = \left(\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n)\right) a^{\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n) - 1}$. Hence, by the Implicit Function Theorem, there exists $s^* = \sigma(a)$ (referred to also as $s^*(a)$ for convenience), which increases in a if and only if $\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n) > 0$. □

A.2. Proof of Proposition 1

Proof. By integrating the slope in Eq. (11) and adding the initial condition $U[0]$, the equilibrium utility for worker i can be expressed as:

$$U[i] = \int_0^i \left(\frac{1}{\psi}\right) \left(a\left[\,\tilde{i}\,\right]\right)^{\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n)} s'\left[\,\tilde{i}\,\right] d\tilde{i} + U[0]. \tag{A.3}$$

Since the ranking of workers' talent in the upper tail satisfies $s'[i] = B(1-i)^{-\beta-1}$ and the productivity of a firm's job position with ranking j is $a[j] = a_L (1-j)^{-\gamma}$, we can substitute these expressions into Eq. (A.3). In particular, Lemma 1 implies that, if $\theta + \sum_{n=1}^{N}(\kappa_n \alpha_n) > 0$, PAM holds and $j^*(i) = i$, so that we can write that:

$$
\begin{aligned}
U[i] &= \int_0^i \left(\frac{1}{\psi}\right) \left(a_L \left(1-\tilde{i}\right)^{-\gamma}\right)^{\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)} B \left(1-\tilde{i}\right)^{-\beta-1} d\tilde{i} + U[0] \\
&= \left(\frac{1}{\psi}\right) (a_L)^{\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)} B \int_0^i \left(1-\tilde{i}\right)^{-\gamma\left(\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)\right) - \beta - 1} d\tilde{i} + U[0] \\
&= \left(\frac{1}{\psi}\right) \frac{(a_L)^{\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)} B}{\gamma\left(\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)\right) + \beta} \left(\left(\left(\frac{a[i]}{a_L}\right)^{-\frac{1}{\gamma}} \right)^{-\gamma\left(\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)\right) - \beta} - 1 \right) + U[0] \\
&= \frac{B}{\gamma\left(\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n) + \frac{\beta}{\gamma}\right)} \left(\frac{1}{\psi}\right) (a_L)^{\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)} \left(\left(\frac{a[i]}{a_L}\right)^{\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n) + \frac{\beta}{\gamma}} - 1 \right) + U[0].
\end{aligned}
\tag{A.4}
$$

We further require $\gamma(\theta + \sum_{n=1}^{N} (\kappa_n \alpha_n)) + \beta > 0$ to ensure the integral defining equilibrium utility converges. The proposition thus follows given the specified initial condition under Assumption 1. Moreover, substituting Eq. (12) into Eq. (9), and then into $e[i] \equiv \sum_{n=0}^{N} x_n[i]\, c_n\, a[i]^{-\kappa_n}$, we obtain the expression for the equilibrium expenditure.

□

A.3. Proof of Proposition 2

Proof. Under PAM, where firm j is matched with $i^*(j) = j$, its profit can be rewritten as:

$$
\overline{V}(j) = a[j]^{\theta} s[i^*(j)] - a[j]^{-\Sigma_n(\kappa_n \alpha_n)} \left(\int_0^{i(j)} a[j(\tilde{i})]^{\theta + \Sigma_n(\kappa_n \alpha_n)} s'[\tilde{i}] d\tilde{i} + u_L \right). \tag{A.5}
$$

Note that the revenue is unaffected as the matching remains the same. The change in the profit thus depends on the change in the total expenditure to workers. Setting $q = \Sigma_n (\kappa_n \alpha_n)$, we have that:

$$
\frac{\partial \overline{V}(j)}{\partial q} = - \left\{ \int_0^i a[j(\tilde{i})]^{\theta} \left(\frac{a[j(\tilde{i})]}{a[j]} \right)^q \underbrace{\left(\ln \frac{a[j(\tilde{i})]}{a[j]} \right)}_{<0} s'[\tilde{i}] d\tilde{i} + \underbrace{\frac{d}{dq} \left(a[j]^{-q} u_L \right)}_{<0} \right\} > 0. \tag{A.6}
$$

Since $\frac{\partial \overline{V}(j)}{\partial \alpha_n} = \frac{\partial \overline{V}(j)}{\partial q} \kappa_n$, it therefore follows that firm profits increase with α_n iff $\kappa_n > 0$.

For the expenditure, taking the partial derivative of Eq. (19) with respect to α_n, yields:

$$\begin{aligned}\frac{\partial e[i]}{\partial \alpha_n} &= \left(B\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)(a[i])^{\theta+\frac{\beta}{\gamma}}\left(-\frac{\kappa_n}{\left(\gamma\left(\theta+\sum\limits_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}\right)\right)^2}\right)\\ &= \underbrace{\left(\frac{B}{\gamma\left(\theta+\sum\limits_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}\right)}\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)(a[i])^{\theta+\frac{\beta}{\gamma}}}_{e[i]}\left(-\frac{\kappa_n}{\gamma\left(\theta+\sum\limits_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}\right)}\right)\\ &= e[i]\left(-\frac{\kappa_n}{\gamma\left(\theta+\sum\limits_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}\right)}\right).\end{aligned} \tag{A.7}$$

The partial elasticity of firm i's equilibrium expenditure with respect to workers' utility weight for amenity n is:

$$\frac{\partial \ln(e[i])}{\partial \ln(\alpha_n)} = \frac{\partial e[i]}{\partial \alpha_n}\frac{\alpha_n}{e[i]}. \tag{A.8}$$

Substituting Eq. (A.7) into Eq. (A.8) yields Eq. (20). The partial elasticity of firm i's equilibrium expenditure with respect to firms' advantage in the provision of amenity n is computed following the same logic. $\square$

A.4. Proof of Corollary 3

Proof.

$$e(i) = \left(\frac{B}{\gamma\left(\theta+\sum_{n=1}^{N}\kappa_n\alpha_n\right)+\beta}\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)(a(i))^{\theta+\frac{\beta}{\gamma}},$$

$$Q(i) = a(i)^{\theta}s(i)L(i), \qquad s(i) = s_H - \left(\frac{1-i}{\widetilde{B}}\right)^{-\beta}, \qquad \beta < 0.$$

$$\frac{e(i)L(i)}{Q(i)} = \frac{\left(\frac{B}{\gamma\left(\theta+\sum_{n=1}^{N}\kappa_n\alpha_n\right)+\beta}\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)a(i)^{\theta+\frac{\beta}{\gamma}}}{a(i)^{\theta}s(i)}.$$

Canceling powers of $a(i)$,

$$\frac{e(i)L(i)}{Q(i)} = \frac{B}{\gamma\left(\theta + \sum_{n=1}^{N}\kappa_n\alpha_n\right) + \beta}\left(\frac{a(i)}{a_L}\right)^{\frac{\beta}{\gamma}}\frac{1}{s(i)}.$$

Using

$$a(i) = a_L(1-i)^{-\gamma},$$

we obtain

$$\left(\frac{a(i)}{a_L}\right)^{\frac{\beta}{\gamma}} = \left((1-i)^{-\gamma}\right)^{\frac{\beta}{\gamma}} = (1-i)^{-\beta}.$$

Thus,

$$\frac{e(i)L(i)}{Q(i)} = \frac{B}{\gamma\left(\theta + \sum_{n=1}^{N}\kappa_n\alpha_n\right) + \beta}\frac{(1-i)^{-\beta}}{s_H - \left(\frac{1-i}{\widetilde{B}}\right)^{-\beta}}.$$

Finally, since $\beta < 0$, we have that:

$$B = -\beta\left(\frac{1}{\widetilde{B}}\right)^{-\beta} = -\beta\,\widetilde{B}^{\beta} > 0,$$

which implies

$$\left(\frac{1-i}{\widetilde{B}}\right)^{-\beta} = (1-i)^{-\beta}\widetilde{B}^{\beta}.$$

Hence, with $\beta < 0$, it follows that:

$$\frac{e(i)L(i)}{Q(i)} = \frac{-\beta}{\gamma\left(\theta + \sum_{n=1}^{N}\kappa_n\alpha_n\right) + \beta}\frac{\widetilde{B}^{\beta}(1-i)^{-\beta}}{s_H - \widetilde{B}^{\beta}(1-i)^{-\beta}}.$$

□

A.5. Equilibrium Robustness

The derivation of the aforementioned propositions is based on imposing the initial condition $U[0] = \frac{B}{\gamma\left(\theta + \sum\limits_{n=1}^{N}(\kappa_n\alpha_n) + \frac{\beta}{\gamma}\right)}\left(\frac{1}{\psi}\right)(a_L)^{\theta + \sum\limits_{n=1}^{N}(\kappa_n\alpha_n)}$. However, instead of requiring this initial condition, we can alternatively follow Gabaix and Landier (2008) and consider the domain of very large firms.

We can take the limit of Eq. (A.3), and subsequently the limit of Eqs. (9), and (10), as $i \to 1$. It then turns out that workers' equilibrium utilities and compensation bundles, along with firms' equilibrium expenditures and profits, can still be approximately expressed as power functions of the corresponding firm size. They do not have Pareto quantile functions (i.e., it is not suitable to say that in this case), but the empirical equations are the same.

More analytically, the last line of Eq. (A.4) can be rewritten as follows:

$$U[i] = \frac{B}{\gamma\left(\theta + \sum_{n=1}^{N}(\kappa_n\alpha_n) + \frac{\beta}{\gamma}\right)}\left(\frac{1}{\psi}\right)\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\left((a[i])^{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}} - a_L^{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}}\right) + U[0]. \tag{A.9}$$

As in Gabaix and Landier (2008), we consider the domain of very large firms by taking the limit of Eq. (A.9) as $i \to 1$. Then, $(a[i])^{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}}$ becomes very large relative to $a_L^{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}}$ and $U[0]$, so that:

$$U[i] = \frac{B}{\gamma\left(\theta + \sum_{n=1}^{N}(\kappa_n\alpha_n) + \frac{\beta}{\gamma}\right)}\left(\frac{1}{\psi}\right)\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}(a[i])^{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)+\frac{\beta}{\gamma}}. \tag{A.10}$$

which is the same expression as the one in the first line of Eq. (12).

Consequently, when the emphasis is on large publicly traded firms, the empirical analysis can still be implemented using the same set of Eqs. (17) and (16).

B. Other Parameters

For completeness, we additionally provide estimates for the other parameters of our model.

Estimating $1-\alpha_0$ based on profits and wages. We utilize the following proposition to estimate $1-\alpha_0$.

Proposition B.1. *The relationship between profit and wages, scaled by assets, yields that:*

$$\frac{\overline{V}[i]}{(a[i])^{\theta}} = \phi + \tau\,\frac{w[i]}{(a[i])^{\theta}}, \tag{B.1}$$

where $\phi \equiv \begin{cases} 0 \text{ if } \beta > 0 \\ s_H \text{ if } \beta < 0 \end{cases}$ *and* $\tau \equiv \frac{1}{\alpha_0}\left(\frac{\theta + \sum_{n=1}^{N}(\kappa_n\alpha_n)}{\frac{\beta}{\gamma}}\right)$.

First of all, since the wage bill constitutes α_0 of total expenditure, one dollar change in wages means $\frac{1}{\alpha_0}$ dollar increase in the total expenditure. Hence, a lower α_0 (and a higher τ) implies a higher change in profit for given any difference in wages. This is because firms need to spend on other amenities when workers do not care that much about the wage.

Second, notice that whether the ratio $\frac{w[i]}{(a[i])^{\theta}}$ increases or decreases in ranking depends on the sign of β, as follows:

$$\frac{w\left[i\right]}{\left(a\left[i\right]\right)^{\theta}}=\frac{\alpha_0 e[i]}{\left(a\left[i\right]\right)^{\theta}}=\left(\frac{B}{\gamma\left(\theta+\sum_{n=1}^{N}\left(\kappa_n\alpha_n\right)+\frac{\beta}{\gamma}\right)}\left(\frac{1}{a_L}\right)^{\frac{\beta}{\gamma}}\right)\left(a\left[i\right]\right)^{\frac{\beta}{\gamma}}. \tag{B.2}$$

To understand this, recall that, if $\beta>0$, then workers' talent distribution is Pareto, thereby implying that there is a good chance of having highly skilled workers at the top. In this case, $\frac{e[i]}{(a[i])^{\theta}}$ increases with respect to the ranking index i, to capture the notion that the expense ratios of firm positions become higher at the top in order to hire those highly skilled workers. On the other hand, if $\beta<0$, workers' talent distribution has an upper bound of s_H, and thus becomes more alike at the top. In that case, $\frac{e[i]}{(a[i])^{\theta}}$ decreases with respect to i to capture the notion that it is relative inexpensive to hire skilled workers at the top. But in either case, the higher the absolute value of β, the larger the difference in the expenditure ratios across firm positions (i.e., $\frac{d}{di}\ln\left(\frac{e[i]}{(a[i])^{\theta}}\right)=\frac{\beta}{1-i}$).

Together with the definition of τ (which is positive iff $\beta>0$), Eq. (B.1) implies that, regardless of the sign of β: (i) $\frac{\overline{V}[i]}{(a[i])^{\theta}}$ always increases in the ranking i, and (ii) a lower value of α_0 predicts that the distribution of firms' profitability ratios is more dispersed relatively to the distribution of wage expense ratio. The practicality of the above proposition is that, by using data on firms' net income (or income before extraordinary items), wages, assets, and the calibrated value of the parameter θ, we can estimate workers' utility weight on wages (α_0), and thus infer the extent of their non-pecuniary preferences.

We therefore run the regression in Eq. (B.1) to estimate τ, and subsequently solve for α_0, which is given by $\alpha_0=\frac{1}{\tau}\frac{\theta+\sum_{n=1}^{N}(\kappa_n\alpha_n)}{\frac{\beta}{\gamma}}$. If workers value firms' amenities – rather than deriving utility solely from wages – then α_0 should be less than 1. In contrast, the purely pecuniary model imposes $\alpha_0=1$, which can be considered the null hypothesis. In particular, we run:

$$\frac{\frac{Profits_{i,t}}{EMP_{i,t}}}{Assets_{i,t}^{\hat{\theta}}}=\underbrace{\frac{1}{\alpha_0}\left(1+\frac{\sum_{n=1}^{N}\kappa_n\alpha_n}{\theta}\right)}_{\equiv\tilde{\tau}}\frac{\hat{\theta}}{\widehat{\beta/\gamma}}\frac{Wage_{i,t}}{Assets_{i,t}^{\hat{\theta}}}+\eta_t^{\overline{V}/a^{\theta}}+\eta_i^{\overline{V}/a^{\theta}}+\epsilon_{i,t}^{\overline{V}/a^{\theta}}, \tag{B.3}$$

Table B.1
Regressions of a firm's profit-per-worker to adjusted-assets ratio on its adjusted-wage-to-adjusted-assets ratio

This table presents regressions of a firm's profit-per-worker to adjusted-assets ratio on its adjusted-wage-to-adjusted-assets-ratio. The dependent variable is $\frac{Profits_{i,t}}{EMP_{i,t}}/Assets^{\hat{\theta}}_{i,t}$, i.e., firm i's profit-per-worker relative to adjusted assets in year t. In Column 1, it is constructed using firm i's income before extraordinary items (IB). In Column 2, it is constructed using firm i's net income (NI). In all columns, the independent variable is $\frac{\hat{\theta}}{\widehat{\beta/\gamma}}Wage_{i,t}/Assets^{\hat{\theta}}_{i,t}$, i.e., firm i's adjusted-wage-to-adjusted-assets ratio in year t. The table depicts the coefficient estimates and the bootstrapped standard errors clustered by firm and stratified by the number of years in a firm is observed (in parenthesis). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. See Table 3 for a detailed description of the sample.

	(1)	(2)
Dependent variable based on	IB	NI
$\frac{\hat{\theta}}{\widehat{\beta/\gamma}}Wage_{i,t}/Assets^{\hat{\theta}}_{i,t}$	1.778***	1.842***
	(0.579)	(0.587)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	8,436	8,436
R^2	0.247	0.263

where $\frac{\frac{Profits_{i,t}}{EMP_{i,t}}}{Assets^{\theta}_{i,t}}$ is firm i's profit-per-worker relative to its adjusted assets (i.e., the output of its assets) in year t, while $\frac{\hat{\theta}}{\widehat{\beta/\gamma}}\frac{Wage_{i,t}}{Assets^{\hat{\theta}}_{i,t}}$ is firm i's adjusted-wage-to-adjusted-assets ratio in year t. Both variables are generated using the inferred value of θ. As usual, $\eta_t^{\overline{V}/a^{\theta}}$ is year t's fixed effect, $\eta_i^{\overline{V}/a^{\theta}}$ is firm i's fixed effect, and $\epsilon_{i,t}^{\overline{V}/a^{\theta}}$ is the error term.

The estimation results are presented in Table B.1. A firm's profit is measured by income before extraordinary items (IB) in Column 1, and by net income (NI) in Column 2. Regardless of the profit measure that is used to generate the dependent variable, the estimated coefficient of $\frac{\hat{\theta}}{\widehat{\beta/\gamma}}\frac{Wage_{i,t}}{Assets^{\hat{\theta}}_{i,t}}$ is close to 1.8 and statistically significant with a t-statistic greater than 3 (based on bootstrapped standard errors that account for the sampling error in the estimation of θ and $\frac{\beta}{\gamma}$, and the unbalancedness of our panel data).

To visualize this, we plot the residualized ratio of IB-per-worker relative to adjusted assets against the residualized adjusted-wage-to-assets ratio in Subfigure B.1a, and the corresponding plot using NI-per-worker in Subfigure B.1b. Each point in the two plots represents a firm-year observation, with all variables residualized with respect to firm and year fixed effects. The x-axis is grouped into 100 bins to aid visual clarity. In both subfigures, the slope of the fitted line is strongly positive and around 1.8, consistent with the above regression results. Invoking the analytical expression of $\tilde{\tau}$, we then obtain a value of between 0.7990 and 0.818 for α_0.

Estimating π using firm assets. We begin by examining the distribution of firm size, focusing specifically on the Pareto fit of firm assets. Following the approach of Gabaix and Ibragimov (2011), we estimate the shape parameter π by running the following regression:

$$\ln(Assets_{i,t}) = -\pi \ln(IndustryRank_{i,t} - \frac{1}{2}) + \eta^{Assets}_{Industry_i} + \eta^{Assets}_t + \epsilon^{Assets}_{i,t}, \tag{B.4}$$

where $\ln(Assets_{i,t})$ is the natural log of firm i's assets in year t, $\ln(IndustryRank_{i,t} - \frac{1}{2})$ is the natural log of the rank of firm i's assets in its industry in year t minus the 1/2 term, $\eta^{Assets}_{Industry_i,t}$ is firm i's industry fixed effect, η^{Assets}_t is year t's fixed effect, and $\epsilon^{Assets}_{i,t}$ is the error term. Essentially, the parameter π is estimated from the cross-section of firm asset sizes after controlling for industry- and year-specific shocks through fixed effects.

The regression results are presented in Panel A of Table B.2. The estimated coefficient of $\ln(IndustryRank_{i,t} - \frac{1}{2})$ implies that π equals 1.158, which is close to one – the value consistent with Zipf's law for firm size.

Estimating m using firm total number of employees' and assets. Next, we use the functional form of $L(a)$, where we assume that the total number of employees is a concave power function of firms' assets, to estimate m. In particular, we take logs of the functional form of $L(a)$ and obtain the following regression equation:

$$\ln(EMP_{i,t}) = m \ln(Assets_{i,t}) + \eta^{EMP}_t + \eta^{EMP}_{Industry_i} + \epsilon^{EMP}_{i,t}, \tag{B.5}$$

where $\ln(EMP)_{i,t}$ is the natural log of firm i's number of employees in year t, $\ln(Assets_{i,t})$ is the natural log of firm i's assets in year t, $\eta^{EMP}_{Industry_i,t}$ is firm i's industry fixed effect, η^{EMP}_t is year t's fixed effect, and $\epsilon^{EMP}_{i,t}$ is the error term. That is once again, we control for industry- and year-specific shocks.

The output of this regression is presented in Panel B of Table B.2. The coefficient of $\ln(Assets_{i,t})$ is found to be 0.662.

Estimating γ and β from the estimates of the other parameters. Having estimated the values of π and m, we can now estimate the value of γ. In particular, since $\gamma \equiv \pi - m$, it follows that it is equal to 0.496. Furthermore, since we have estimated the value of β/γ, we can subsequently estimate the value of $\beta = \gamma \cdot (\beta/\gamma)$. Specifically, our estimates imply that β is equal to -0.050.

Fig. B.1. Scatterplots of the firm's profit-per-worker to adjusted-assets ratio against the firm's adjusted wage-to-adjusted-assets ratio. Each subfigure plots the relationship between the residualized profit-per-worker to adjusted-assets ratio (based on either income before extraordinary items or net income) and the residualized adjusted-wage-to-adjusted-assets ratio. Residuals are obtained by regressing both variables on firm and year fixed effects. The y-axis shows the residualized value of $\frac{Profits_{i,t}}{EMP_{i,t}}/Assets^{\hat{\theta}}_{i,t}$, i.e., firm i's profit per worker relative to adjusted assets in year t. In Subfigure B.1a, profits are measured using income before extraordinary items (IB). In Subfigure B.1b, profits are measured using net income (NI). The x-axis shows 100 bins of the residualized $\frac{\hat{\theta}}{\widehat{\beta/\gamma}}Wage_{i,t}/Assets^{\hat{\theta}}_{i,t}$, i.e., firm i's adjusted-wage-to-adjusted-assets ratio in year t.

(a) IB-per-worker relative to adjusted assets against adjusted-wage-to-adjusted-assets ratio

(b) NI-per-worker relative to adjusted assets against adjusted-wage-to-adjusted-assets ratio

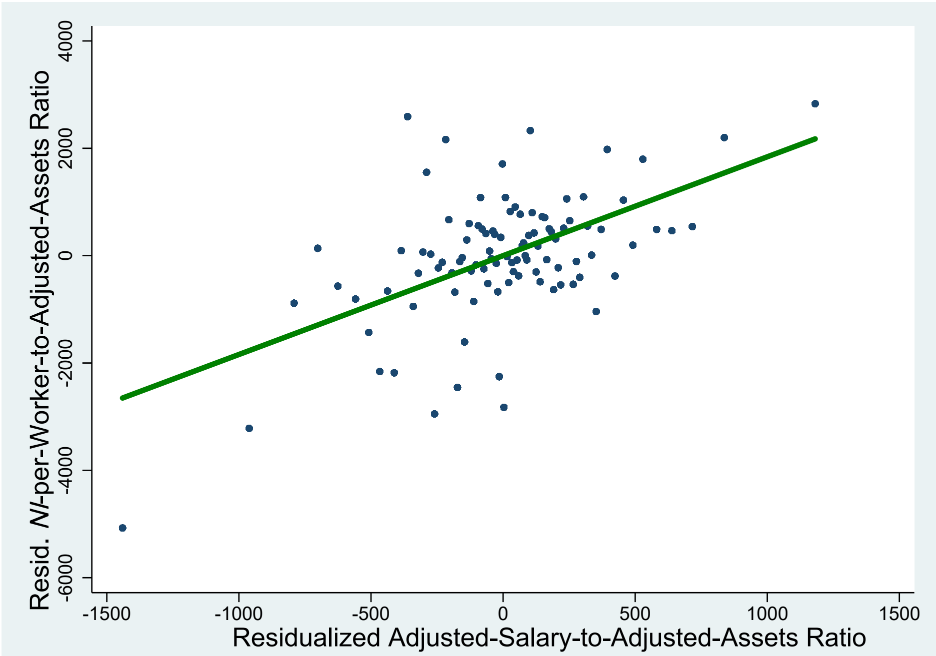

Table B.2

The Pareto fit of firms' assets and the regression implied by the functional form of $L(a)$

This table presents the regressions used to estimate the model's auxiliary parameters. Panel A presents the Pareto fit of firm assets. The dependent variable is $\ln(Assets_{i,t})$, i.e., the log of firm i's assets in year t. The independent variable is $\ln(IndustryRank_{i,t} - \frac{1}{2})$, i.e., the log of the rank of firm i's assets in its industry in year t minus the 1/2 term. Panel B presents the regression implied by the functional form of $L(a)$. The dependent variable is $\ln(EMP_{i,t})$, i.e., the log of firm i's number of employees in year t. The independent variable is $\ln(Assets_{i,t})$, i.e., the log of firm i's assets in year t. The industries are the 11 two-digit GICS industries. The table depicts the coefficient estimates and the two-way clustered standard errors at the firm and year level (in parenthesis). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. See Table 3 for a detailed description of the sample.

Panel A: Pareto fit of firm assets	
$\ln(IndustryRank - \frac{1}{2})$	-1.158***
	(0.034)
Industry FE	Yes
Year FE	Yes
Observations	8,534
R^2	0.837
Panel B: Regression implied by the functional form of $L(a)$	
$\ln(Assets)$	0.662***
	(0.033)
Industry FE	Yes
Year FE	Yes
Observations	8,459
R^2	0.583