NBER WORKING PAPER SERIES

NON-WAGE AMENITIES

Alexandre Mas

Working Paper 33643 http://www.nber.org/papers/w33643

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

I thank Martin Beraja, Sydnee Caldwell, Patrick Kline, Andreu Mas-Colell, Bailey Palmer, Sebastian Puerta, Raffaele Saggio, Reed Walker and editors Christian Dustmann and Thomas Lemieux for helpful comments. I am grateful to Kyle Neary for excellent research assistance. ROCKWOOL Foundation Berlin (RFBerlin) provided funding for the Handbook conference. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

 \bigcirc 2025 by Alexandre Mas. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including \bigcirc notice, is given to the source.

Non-Wage Amenities Alexandre Mas NBER Working Paper No. 33643 April 2025 JEL No. J0, J3, J30, J31, J32, J33

ABSTRACT

This chapter reviews the analysis of non-wage amenities in the workplace. The competitive model is the point of departure, but the emphasis is on models of imperfect competition that have greater empirical relevance. In addition to the traditional hedonic model for estimating preferences over job characteristics, I review revealed-preference methods that are better suited when there are deviations from perfect competition.

Alexandre Mas University of California, Berkeley Haas School of Business 2220 Piedmont Ave Berkeley, CA 94720 and NBER amas@berkeley.edu

Contents

1	Intr	oduction	3
2	Am	enity Determination Under Perfect Competition	6
	2.1	Benchmark Model of Amenity Determination	6
		2.1.1 Cobb-Douglas Calibration for the Marginal Value of Health	
		Benefits	9
	2.2	Which Objects Do We Want to Measure?	10
		2.2.1 Demand System Approach	10
	2.3	Compensating Differentials	11
	2.4	Public Good Amenities	14
3	Ben	chmark Comparative Statics and Other Sources of Amenity Variation	16
	3.1	Variation in Cost of Amenity Provision	16
	3.2	Variation in Preferences Over Amenities	18
	3.3	Variation in Market Utility	18
	3.4	Other Mechanisms	20
		3.4.1 Productivity	20
		3.4.2 Dynamic Considerations	21
		3.4.3 Unions	21
		3.4.4 Inside the Firm Heterogeneity and Spillovers	22
4	Two	-Sided Heterogeneity and Sorting: The Rosen Model	24
	4.1	A Worked Out Case with a Continuous Amenity	25
		4.1.1 Long-run Equilibrium	28
	4.2	Discrete Case	30
	4.3	Evidence on Sorting	32
5	Hed	onic Models: Estimation	33
	5.1	Hedonic Models: Assumptions	34
	5.2	Compensating Differentials in Two Settings	37
		5.2.1 Working From Home as a Job Attribute	37
		5.2.2 Injury and Fatality Risk	40
6	Mar	ket-level Incidence and Mandates	46
7	The	Analysis of Non-Wage Amenities Under Imperfect Competition	49
	7.1	Amenity Determination with Search Frictions	50

		7.1.1 Static Exposition as a Monopsony Problem	51
	7.2	Horizontal Amenities and Thinness in the Distribution of Preferences .	52
8	Rev	ealed Preferences Approaches	55
	8.1	Discrete Choice Experiments	56
		8.1.1 Survey-Based Discrete Choice Experiments	58
	8.2	MWTP Estimation with Non-Experimental Offer Data	60
	8.3	Using Worker Mobility to estimate MWTP for Job Characteristics	65
	8.4	Using Flows to Estimate Firm Value	68
		8.4.1 Application to Collective Bargaining Agreements	71
		8.4.2 Application to Estimation of Work Hour Constraints	72
	8.5	Equilibrium Models	73
9	Con	clusion	75

1 Introduction

The value of a job is the combination of its attributes, some of which are monetary, and others non-monetary. Non-wage work amenities are an important part of this overall value. Traditional benefits like health insurance, retirement and paid leave alone represent approximately 30 percent of total worker compensation (BLS 2024) in the United States, on average. However, benefits are only one dimension of the value of a work-place. Scheduling practices, family leave, and telecommuting policies have all been shown to have relevance to many workers (Mas and Pallais 2017; Wiswall and Zafar 2018; Maestas et al. 2023). Intangible amenities, like dignity at work and whether the workplace has a hostile climate, weigh heavily in worker preferences (Dube, Naidu, and Reich 2022; Collis and Van Effenterre 2025). In the employer-employee matched Longitudinal Employer-Household Dynamics (LEHD) data, forty percent of employment-to-employment moves are to lower paying firms (Sorkin 2018), suggesting that motivations beyond wages drive many job transitions.

A second class of facts concern differences in availability or take-up of non-wage amenities by worker characteristics.¹ The body of evidence consistently point to the conclusion that valuable non-wage amenities and better working conditions are disproportionately offered at higher parts of the earnings distribution. In 2014, the benefits share of total compensation for civilians was over twice as large at the 80th percentile of the wage distribution than at the 20th percentile (Monaco and Pierce 2015). Workers with higher education have better working conditions (Maestas et al. 2023). Jobs with better amenities in one dimension have better amenities in other dimensions (Sockin 2022). Jobs with better amenities, on average, tend to also have higher wages (Pierce 2001; Roussille and Scuderi 2023; Sockin 2022; Caldwell, Haegele, and Heining 2024). But there are nuances, for example evidence that in some cases less productive workers select to have amenities, like working from home (WFH), when given the option (Emanuel and Harrington 2024). There are also documented differences in availability and consumption of non-wage amenities by gender, age, and other demographic characteristics. For example, Goldin, Pekkala Kerr, and Olivetti (2020) document that in the U.S. there is a positive correlation of firms that offer parental paid leave and the share of the workforce that is of childbearing age. In Brazil, women are more likely to hold jobs with better non-wage amenities (Morchio and Moser 2024). Female college graduates are more likely to be in jobs with greater work flexibility and job stability (Wiswall and Zafar 2018).

¹In this Chapter, I will use the term non-wage amenity broadly to refer to any job characteristic that is not direct compensation.

What explains the high levels and variation of non-wage amenities across firms and workers? Labor economists have confronted the growing evidence base, building frameworks to understand their effects on overall welfare, the wage structure, how workers value these amenities, and how firms produce them. In this Chapter, I discuss these frameworks—both theoretical and econometric—and the empirical findings. I begin with a benchmark representative agent model of amenity determination under perfect competition. While for historical reasons this topic is often introduced with the more complex model of two-sided heterogeneity from Rosen (1986) and Rosen (1974), which I will also review, many of its key predictions can be understood from this simpler and more tractable framework. I use this framework to guide the analysis of the determinants of non-wage amenities. Variation in amenities can be driven by costs, preferences, income, or other factors like unions, productivity investments, or regulation.

I also use the benchmark model to discuss what has been the traditional focus on amenities: the theory of equalizing differences or compensating differentials, under which monetary and non-monetary factors balance a job's overall attractiveness, and hedonic models, which examine how wages adjust to for varying job characteristics, reflecting worker preferences and firms' costs. There are several reasons why the theory of equalizing differences and compensating differentials is a pillar of the labor economics canon. The first is that it is a sharp test of competitive markets. By the definition of a competitive equilibrium, in a static equilibrium workers should not improve their situation by changing firms. When there is an undesirable job characteristic, wages should be higher to maintain this condition. This is, at least in principle, a straightforward prediction to test and one that is at the core of our understanding of market structure. A second reason for the interest in compensating differentials is that it is a mechanism explaining variation in wages. Why do wages vary for the same worker across different firms, as has been documented extensively (Abowd, Kramarz, and Margolis 1999 (AKM))? It is a candidate to explain, without resorting to market inefficiency, why the law of one price seemingly does not hold in the labor market-this variation may just reflect amenity differences. In the same vein, compensating differentials are closely aligned to incidence questions that are also germane to the wage structure-for example, how rising health care costs affect wages and inequality, or how benefit mandates, like paid family leave, affect the gender pay gap. The concept is also useful for understanding spatial differences in wages, following the seminal paper of Roback (1982) and the large literature that has followed. A third reason for the interest in compensating differentials is its application to determine implicit prices, that is, marginal willingness to pay (MWTP) for non-market work characteristics. MWTP

can be used to evaluate policies affecting the work environment. For example, if we are interested in the welfare implications of regulating irregular scheduling, it would be useful to know workers' willingness to pay (WTP) to avoid irregular scheduling and to weigh that against the cost of the policy.

In this chapter, I follow my interpretation of the direction of the literature by broadening the inquiry beyond compensating differentials and hedonics to the more general question on the determinants and implications of non-wage amenities. One motivation for this more expansive approach is that hedonic approaches to estimate compensating differentials have been generally incapable producing reliable estimates of preferences over amenities. There are many reasons, but an important one is that if the perfect competition test fails—and by my read it will—the utility of subsequent theoretical applications premised on a competitive model diminishes. I subsequently review approaches for analyzing and incorporating non-wage amenities that assume or embrace models of imperfect competition. These models can include rent-sharing, efficiency wages, monopsony, and search frictions. These models share the prediction that, once we deviate from perfect competition, there is variation in job value for a given worker, and this variation will be in conflict with the law of one price. This variation in job value turns out to be important quantitatively, and breaks the relationship between wages and amenities that we might otherwise observe in a compensating differentials framework. My discussion also includes evaluating the now commonly employed approach of explaining markdown in wages by way of horizontal firm differentiation.

I then outline different revealed preferences approaches that look directly to the choice margin to estimate preference parameters. These include experimental approaches, as in Mas and Pallais (2017), discrete choice experiments embedded in surveys, methods that use non-experimental offer data or information on the job separation margin, applications from the computer science literature to use employment flows to estimate the value of firms, as in Sorkin (2018), and structural equilibrium models that infer the aggregate importance of non-wage amenities from observed moments. These revealed preference methods have been increasingly used because they do not rely on assumptions of perfect competition. Indeed, they are premised on imperfectly competitive markets. Advances in these approaches have expanded the production possibilities frontier in this broad area of inquiry.

Traditionally, when non-wage amenities are taught, students encounter the "mixed record" of hedonic compensating differential estimates, and the topic often ends more or less there on a pessimistic note. My aim in this Chapter is to emphasize possibilities. And I believe this optimism is warranted. There is now a large and mature literature on non-wage amenities that has broadly that has influenced all areas of economics beyond

labor, including macroeconomics, public finance, theory, and econometrics. It is also an area of inquiry that cuts across methodological silos, including price and search theory, randomized control trials, survey experiments, and structural estimation. All of these different approaches have led to numerous substantive findings on how workers and firms value amenities, and on the structure of markets. Beyond the academic literature, this is an area with a close connection to applications, ranging from the increasing use of conjoint analysis by companies to determine employee benefit packages, to the estimation of the value of a statistical life (VSL), a key input in cost-benefit analysis of regulation.

2 Amenity Determination Under Perfect Competition

I begin by describing models of amenity determination in a perfectly competitive market, starting from the problem of the firm as the consumer, which I call the benchmark model, followed by the influential Rosen (1974) and Rosen (1986) models that allow for worker and firm heterogeneity and sorting.

2.1 Benchmark Model of Amenity Determination

Why do firms provide non-wage amenities instead of higher pay? A plausible answer is that the value of amenities to workers exceeds the average cost of amenity provision to firms. I formalize this in a representative agent model for optimal amenity choices. Variants of this model have been employed by Woodbury and Huang (1991), Sorkin (2018) and Dube, Naidu, and Reich (2022).

Consider a firm choosing a wage *w* and *n* real-valued amenities $\{a_1, a_2, ..., a_n\}$ for a representative worker. Denote the tax rate on wages *t* and net-of-tax wages $z \equiv (1-t)w$. Amenities are untaxed. Let the worker's utility be $V(z, a_1, ..., a_n)$, where $\frac{\partial^2 V}{\partial a_k^2} \leq 0$ for all $k \in \{1, ..., n\}$. The total value of a job for a given level of skill *s*, denoted $\overline{V}(s)$, is determined in the competitive market and taken as given by the firm. The firm faces per-unit price c_k for each unit of amenity a_k , and pays a wage *w*.

The firm problem is to minimize expenditure subject to the worker reaching at least its target utility $\overline{V}(s)$:

$$C\left(\overline{V}(s), c_1, \dots, c_n, 1/(1-t)\right) = \min_{w, a_1, \dots, a_n} \left[w + \sum_{k=1}^n c_k a_k\right],$$

subject to $V(z, a_1, \dots, a_n) \ge \overline{V}(s),$ (1)

where $C(\overline{V}(s), c_1, \ldots, c_n, 1/(1-t))$ is the expenditure function that represents the minimal expenditure necessary for the firm to meet the worker's utility, given amenity and wage costs.

This expenditure minimization problem, or the equivalent primal problem of maximizing utility subject to a total cost level, is a fair representation of how many employers actually set benefits. For example, Deloitte advises clients to "identify the ... combination of [wage-amenity] offerings and levels that employees will most prefer at... any given cost" (DeBellis 2018). Formally, it is exactly the consumer problem, where the firm is the consumer and c_k are the prices. In this problem, the after-tax wage is an amenity quantity, like any other, and its price to the firm is $\frac{1}{1-t}$. While the firm is the consumer in this problem, the only relevant preferences are the worker's. In this model, if my employer did not provide me with the extra ergonomic chair and instead opted for a less comfortable one that costs \$200 less, it is because I, the representative worker, am not willing to pay the extra \$200 for it at the margin. The price a worker pays for the ergonomic chair is in the form of some combination of a lower wage and reductions in other amenities. (This example makes clear that we will eventually need to eventually allow for heterogeneity of preferences within the firm.)

Before proceeding it is worth asking, what are these amenities in the firm's objective function? Here, they are anything the worker values in the context of work, that the firm can provide at some cost. This can be fringe benefits, ergonomics, or respectful managers. Average work hours, which has traditionally been analyzed with models of labor supply, can be best thought of as a firm amenity, as emphasized by Lachowska et al. (2023b).² Environmental, social, and governance (ESG) practices are not directly oriented to workers, but tend to be valued and can therefore help in recruitment (Colonnelli et al. 2023). The amenities in the firm's problem include those that have consumption value and those that affect the disutility of work alone. Wages, fringe benefits, or whether the work is fun or inherently meaningful all have consumption value. The provision of personal protective equipment only serves to reduce the disutility of work and does not have consumption value outside of the workplace. In optimization problem (1), we can think of the job value $V(\cdot)$ as representing overall utility relative to not working, so all amenities belong in the utility function.

The assumption of diminishing marginal utility in amenity consumption and constant costs of amenity provision allows us to obtain an interior solution to problem (1).³

²For the purposes of the amenity determination problem, I only consider amenities that are costly to provide at the margin, because a free amenity that workers value should be at a corner solution.

³Giving firms convex costs of amenity provision and specifying utility either as concave in amenities or linear in the amenities' cash value to the worker would also allow for interior solutions. If amenities and costs both enter linearly, variation in costs leads to corner solutions: the amenity is fully provided if





From the first-order conditions we obtain:

$$\frac{c_k}{1/(1-t)} = \frac{\partial V/\partial a_k}{\partial V/\partial z} = \text{MRS}_{a_k,z}.$$

This expression says that the firm produces the amenity until the marginal rate of substitution (MRS) of the amenity with respect to the net-of-tax wage—the MWTP for the amenity—is equal to the ratio of the price of a_k (the c_k term) and the price to the firm of after-tax dollars (1/(1-t)). The optimality condition for the single amenity problem is shown in Figure 1 for a firm targeting utility $\overline{V}(s_1)$.⁴ The optimal bundle is achieved at tangency point (point a) of the indifference curve for $\overline{V}(s_1)$ and the firm's isocost curve. The intercept is the expenditure function corresponding to $\overline{V}(s_1)$, the amenity price to the firm c, and the tax rate t. From the expenditure function it is easy to appreciate that the worker derives surplus from the amenity. Because it is lower than the indifference curve on the chosen bundle, a worker's utility will be strictly lower if the firm pays them their compensation expenditure in cash. That is, a worker is willing to pay more for the full amount of the equilibrium amenity than the wage cut they pay to obtain it.

WTP > c, and not provided if WTP < c. With corner solutions we can model amenities as discrete (fully turned on or off). The fact that many amenities are not at a corner solution suggests that concavity in utility and/or convexity in costs is relevant in these cases.

⁴While I will show the single amenity case graphically, I maintain the multi-amenity model for the main analyses because in certain situations a single amenity is *with* loss of generality.

2.1.1 Cobb-Douglas Calibration for the Marginal Value of Health Benefits

To ground the representative agent benchmark with data, we can briefly assume a Cobb-Douglas, in which case, from the first-order condition the MRS between amenity k and wage can then be expressed in terms of their expenditure shares according to:

$$\mathrm{MRS}_{a_k,z} = \frac{s_k}{s_w} \frac{w^*}{a_k^*}.$$
 (2)

In this expression w^* and a_k^* are optimal wage and amenity quantities, and s_k and s_w are expenditure shares for amenity k and the wage w relative to total compensation. Directionally, this expression implies that amenities that are more highly valued at the margin have higher expenditure share.

If we employ this (possibly) heroic Cobb-Douglas parametrization, we can compute valuations from observed aggregate data. Consider the actuarial value (AV) of employer-sponsored health insurance, which is defined as the average percentage of total healthcare expenses that a plan covers. We can compute the MWTP for the AV by substituting observed empirical moments into expression (2). I use 2022 Census median earnings (Guzman and Kollar 2023) for w^* (=\$48,000), the Affordable Care Act (ACA) Silver health plan AV (Levitt and Claxton 2011) for a_k^* (=70%), the employer expenditure share of health benefits from the 2022 Bureau of Labor Statistics (BLS) Employer Costs for Employee Compensation Survey for private industry for s_k (=0.07), and the expenditure share of wages from this BLS survey for s_w (=0.70).

Substituting these values, expression (2) implies a MWTP for a one percentage point increase in the AV of $MRS_{a_k,z} = \frac{0.07}{0.7} \frac{$48,000}{70} \implies MRS_{a_k,z} \approx 69 . If the theory is correct, the marginal benefit corresponding to a one percentage point change in the AV should be approximately \$69. We can validate this estimate by comparing it to the cash value of a one percentage point increase in employer contributions to average individual health care expenditures. In the 2022 Medical Expenditure Panel Survey, the average annual medical expenditure for non-institutionalized civilians was \$6,765 (AHRQ 2025). A one percentage point increase in the AV of insurance corresponds to a \$68 cash benefit to the worker, assuming no behavioral changes. Pretty close! Note that the comparison is not circular since health insurance is just one of a number of benefits, so its expenditure share in the Cobb-Douglas first-order condition is not pinned down by the level of medical expenditures. Moreover, the model-based MWTP depends on the average AV of 70% which does not factor directly into the average medical expenditure figure. Had the AV been, say, 20%, we would obtain a much larger MWTP.

2.2 Which Objects Do We Want to Measure?

The first-order condition makes evident that knowledge of the equilibrium MWTP (the MRS) for an amenity, even if it is derived from an experiment, is not necessarily informative about worker preferences globally because it is pinned down by the firm's marginal cost. If this marginal cost is low, we will read a low MWTP, and if it is high, the opposite. Assuming that the competitive assumptions hold, it is just the MWTP at a particular point. In fact, under strict perfect competition assumptions, if we know the marginal cost, we do not even need to run the experiment. The marginal cost is the MWTP.

Consider an experiment that lowers the price of an amenity to the firm. This could be a safety inspection with fines that are levied if the workplace violations are not addressed, thus lowering the price of safety. From this, we obtain three estimates: the fine Δc_k , the change in the negative of injuries (i.e., safety) Δa_k , and the change in the wage Δw . The traditional compensating differentials approach uses the ratio $\frac{\Delta w}{\Delta a_k}$ as the MWTP, which under a host of assumptions, that I will discuss, is valid. Alternatively, we can consider Δa_k and Δc_k as movement along the amenity demand function. Note that $\frac{\Delta c_k}{\Delta a_k}$ is an estimate of the marginal cost of a_k to the firm, and since the firm adjusts its safety investment until the marginal cost of providing additional safety equals its marginal benefit, $\frac{\Delta c_k}{\Delta a_k}$ is the firm's MWTP for safety. Since in the cost minimization problem the constraint is worker utility, this MWTP corresponds to worker preferences. If the experiment shows that the marginal cost does not equal a known MWTP, then we conclude that there is a market failure.

2.2.1 Demand System Approach

More useful than a single MWTP value would be MWTP over a range of amenity values, in other words the inverse demand curve for amenities. This is, in part, the appeal of revealed preference approaches that can directly estimate the MWTP distribution, such as in Mas and Pallais (2017). If data permit, it is potentially appealing to directly estimate demand for a_k as a function of c_k . Fringe benefits are a case where we can plausibly define the relevant prices of provision so as to directly calculate demand functions, and some work has been done in this area, for example Woodbury (1983), Woodbury and Huang (1991), Woodbury and Hamermesh (1992) and Gruber and Lettau (2004).

The constant expenditure assumption of the Cobb-Douglas we considered above can be relaxed by using a more flexible utility function, like a translog, or employing the Almost-Ideal Demand System (AIDS) of Deaton and Muellbauer (1980). Employing the latter leads to a system of demand equations relating firm expenditure shares on amenities to the cost of the amenities. For amenities (k = 1, ..., n) we have:

$$s_k = \alpha_k + \frac{w}{(1-t)} + \sum_{l=1}^n \gamma_{kl} \ln c_l + \beta_k \ln \left(\frac{C}{P}\right),$$

where *C* is total expenditures and *P* is the price aggregator which is a function of amenity costs and the tax rate. Adding error terms, this system can be estimated, but has high data and identification requirements that are usually infeasible for most applications. To identify the utility parameters we require data on costs or quantities that vary by firm, exogeneous cost-shifters, and an instrument for total compensation expenditure. Woodbury and Huang (1991) estimate a model like this one at the establishment level by assuming that the cost of all benefits other than health are constant, approximating the cost of health benefits using a hedonic model of the price of health insurance, and instrumenting total expenditures $\frac{C}{P}$ with after-tax expenditures, using the wedge that comes from the tax deductibility of health benefits. Woodbury and Hamermesh (1992) use an AIDS model to estimate fringe benefit expenditures for university faculty and conclude that the Tax Reform Act of 1986, which lowered marginal tax rates, significantly reduced demand for fringe benefits.

The difficulty with applying this approach more broadly is that prices c_k are usually unobserved. One of the promises of the compensating differentials approach that I will discuss (which, ultimately requires considerable and strong assumptions) is that we can use the observed quantities of one amenity (wages) at a known or normalized price (=1, ignoring taxes) to learn about the unobserved prices of the other amenities, and in turn the MWTP.

2.3 Compensating Differentials

Compensating differentials arise when jobs vary in their non-wage attributes and firms adjust wages to attract and retain workers given these differences. One reason to study compensating differentials is to explain differences in wages across firms, industries, occupations, and regions. This is an inherently interesting question and has long been a central motivation for modeling how wages vary with job amenities in equilibrium, beginning with Adam Smith's original observation that "the wages of labour vary with the ease or hardship, the cleanliness or dirtiness, the honourableness or dishonourableness of the employment" (Smith 1776 [2008]). For this "reduced form" purpose, there is no need to map the relationship in wages and amenities to particular model primitives.

Another reason for estimating compensating differentials is to learn about worker or firm preferences. Due in part to the historical (and I believe misplaced) reluctance to using stated preference surveys, and also the limited data availability on amenity prices, there have been many attempts to estimate compensating differentials models as a way to use the market signal to infer the value of amenities. This approach has featured particularly prominently in the estimation and interpretation of the VSL. My discussion in this section is primarily on this more structural motivation.

The idea of the compensating differentials approach to estimating MWTP is that if employers target a market reservation value, any change in the cost of an amenity will translate into an offsetting wage—an equalizing difference—to keep utility constant. This change in the wage is the worker's MWTP for the amenity change. Within the language of the model, we are inferring the MWTP for amenity k from the change in quantity of amenity i, where the price of the amenity i is known. In this case, amenity iis the wage.

Along the indifference surface, and holding all other a_{-k} fixed, $V(w, a_1, ..., a_n) = \overline{V}(s)$, we have $dV = 0 \implies V_w dw + V_{a_k} da_k = 0$ which gives

$$\left.\frac{dw}{da_k}\right|_{V=\bar{V}(s)} = -\frac{V_{a_k}}{V_w} = -\operatorname{MRS}_{a_k,w} = -c_k.$$

This expression says that changes in the amenity (or amenity cost) will be capitalized in the wage. This capitalization, under the maintained assumptions, is the worker's MWTP for the amenity, which is equal to the known marginal cost of provision, at the optimum.⁵

To illustrate this approach we go back to the case of the health benefit where D(a) is the MWTP at *a* dollars of benefits and (1-t) is the effective marginal cost of their provision. Suppose there are no other amenities so the firm can only choose between the mix of wages and health benefits. If we have a policy experiment that slightly changes the tax rate for workers at a single firm, the implied instrumental variables (IV) estimate on the wage, under perfect competition, will be $\frac{dw}{dt} / \frac{da}{dt} \Big|_{V=\bar{V}} = -(1-t)$. This scenario would, in principle, serve as a test of competitive markets. Since we know *ex ante* the marginal cost is the net-of-tax rate, the test is whether a change in the tax results in a reduction in the wage equal to the net-of-tax rate (assuming a zero marginal cost to the firm of provision). Crucially, for the compensating differential to equal the MWTP in this example, it must be that it holds the reservation job value constant. This means that a policy shock affecting the broader market would not necessarily identify the MWTP because it would change the reservation job value. Any change in the wage in that case

⁵If the amenity is discrete, or we are at a corner solution, the interpretation of differences in wages is the worker's constant valuation to go from no provision to full provision. Any variation in this measure across time, space or markets will reflect differences in worker valuation over the amenity.

would be some combination of a shift in supply and demand.

Multi-Margin Optimization In discussing how the wage changes in response to either a change in the value of an amenity or its cost, I have been careful to say that it is assuming all other amenities are held constant. However, that assumption is internally inconsistent with the expenditure minimization problem, which basically guarantees that other amenities are not constant. This application of Le Chatelier's principle in the context of amenity determination is a fundamental problem for the interpretation of the compensating differential as a measure of the MWTP, and it is not solved even if we have experimental variation on amenity production costs. The problem stems from the fact that amenities are bound by the requirement of maintaining the competitive utility target, and a change to one amenity puts all the amenities up for grabs through re-optimization. In words: suppose we change the safety practices at a plant and reduce injuries. In turn, management expenditure minimizes and finds it optimal to change the shift schedules, which reduces costs but which workers dislike. The net result is no change in job utility and no change in the wage, despite an improvement in safety which workers value. The experimentally estimated compensating differential will be 0.

More generally, suppose we have an experiment that varies the cost c_k of amenity a_k . Then the IV estimate of the effect of a_k on w^* is $\beta_{IV} = (\partial w^* / \partial c_k) / (\partial a_k^* / \partial c_k)$. To derive this quantity, we totally differentiate the firm's expenditure $C(\cdot)$ at the optimal wage and amenity values

$$rac{\partial C}{\partial c_k} = rac{\partial w^*}{\partial c_k} + a_k^* + \sum_{j=1}^n c_j rac{\partial a_j^*}{\partial c_k}.$$

From Shephard's Lemma, $\frac{\partial C}{\partial c_k} = a_k^*$, implying

$$\frac{\partial w^*}{\partial c_k} = -\sum_{j=1}^n c_j \frac{\partial a_j^*}{\partial c_k}.$$

$$\beta_{IV} = -\frac{\sum_{j=1}^n c_j \left(\partial a_j^* / \partial c_k\right)}{\partial a_k^* / \partial c_k} \neq -c_k.$$
 (3)

Therefore, the compensating differential does not equal the MWTP. The change in the wage is just whatever change is required to equalize utility after all of the other amenities have adjusted. In this experiment, we would have to be able to measure the entire cost structure in order to back out the MWTP. One rejoinder might be that the confounding cross-terms are second-order. While true, the compensating differential itself is second-order (in the sense that it is a cross-derivative of the Hessian of the expenditure function). This issue is less serious for the case of demand estimation of the compensated demand (how does quantity of amenity a_k vary with c_k) because the quantity of amenity k is a first-order response to its own price.

There are two special cases where the compensating differential with respect to an amenity reflects the MWTP. The first is if there is only one amenity. In that case equation (3) reduces to -c, which is the MWTP. The second case is if all of the crosseffects are zero. Assuming other amenities do not vary may be a reasonable assumption in some situations, like for exogenous place amenities. However, it is hard to tell a story about why a firm would necessarily adjust wages and not, say, work schedules in response to a new safety requirement.

Empirically there are indicators that this an important concern. Clemens (2021) makes this point with respect to multiple margin adjustments for minimum wage increases.⁶ Lee and Taylor (2019) find that as much as 40 percent of the compensating differential from safer workplaces following a safety inspections is through a reduction in fringe benefits rather than the wage. Mas and Pallais (2020) document how attractive work attributes, like flexible scheduling, are often combined with negative attributes, like high-stress responsibilities. These combinations may be optimal in job design, in which case a change in one job attribute would change other complementary attributes.

Given the difficulty in interpreting compensating differentials as a MWTP, we can think of the wage response to a cost shock as a pass-through, which is interesting in its own right, even if it does not have any particular link to worker preferences.

2.4 Public Good Amenities

Because we have considered the optimal amenity mix for a representative worker, we have not had to think further about whether the amenities are private goods, like cash benefits, or public goods that are non-rivalrous in consumption, like improvements in safety practices, better scheduling management, attractiveness of the physical plant or working conditions more generally. Because the public good amenity does not scale with employment, the distinction is relevant once we introduce the employment margin. Consider the following model: a firm maximizes profits $\pi = pF(L) - wL - ca$ subject to the labor market equilibrium $V(a, w) = \overline{V}$, where p is the revenue per worker, L is the number of workers employed, and $F(\cdot)$ is the production function. In equilibrium, $F'(L) = w(a^*; \overline{V})$ where $w(a^*; \overline{V})$ is implicitly defined by $V(a^*, w) = \overline{V}$. For a

⁶The test of whether the minimum wage affects non-wage amenities is itself a test of the competitive equalizing differences model.

private good, the amenity optimality condition (ignoring taxes) is $MRS_{a,w} = c$ for every worker. However, by the Samuelson condition, with the public goods amenity we require $MRS_{a,w} \times L^*(a,w) = c$, where $L^*(a,w)$ is the equilibrium number of workers hired as a function of the amenity level and the wage. From this expression we can see that the isoprofit curve with respect to the wage *w* and amenity *a* (which incorporates employment) is non-linear and exhibits convexity. Its slope is $\frac{c}{L^*}$, which is a function of *a* by way of L^* . For low values of *a*, labor demand is low because $w(a^*; \overline{V})$ is high (to maintain market utility \overline{V}) so every worker hired incurs a high wage cost. Since L^* is low, every incremental unit of *a* has to be offset by large reductions in wages to maintain constant profits, resulting in a steep negative portion of the isoprofit. For higher values of *a*, L^* will be larger, and when there are more workers the firm only needs to cut the wage by small amounts to cover the incremental cost of the amenity to maintain constant profits, since the wage cuts are summed across all workers.⁷

Because the isoprofit curve has convexity, corner solutions are more likely than in the private amenities case. Figure 2 shows one scenario. In Panel A, the solid curve is the isoprofit curve that maximizes profits subject to achieving factor market utility level $V = \overline{V}$. As drawn, the worker's MWTP for *a* exceeds the firm's effective marginal cost for the amenity of c/L^* over the entire range, and the firm sets *a* to its maximum value \overline{a} . The flattening of the isoprofit makes the corner case more likely. Panel B shows a case where there is an interior solution. Because of the flattening isoprofit, an interior solution requires rapidly diminishing marginal utility from the amenity at a fixed utility level. In the extreme case, Leontief preferences guarantee an interior solution.

One take-away from this model is that large firms should have better working conditions. Fixed-costs on the provision of private amenities will have the same implication. The evidence on whether larger firms have better amenities and working conditions depends on the dimension of the amenity but, on balance, there is a positive relationship. The evidence is strong that larger firms have fewer workplace injuries. Larger firms tend to have a wider range of, and more generous benefits. They are also more likely to have full-time positions, something that is highly valued by workers (Dube, Naidu, and Reich 2022; Lachowska et al. 2023b). The size/amenity relationship accords with the findings in the structural model of Lamadon, Mogstad, and Setzler (2022), though we do not require imperfect competition to generate that feature. A fact that challenges the competitive explanation for the amenity-size relationship is the firm-size wage premium. Contrary to the prediction that firms with better amenities scale because they

⁷While this problem has a similar flavor as increasing returns to scale, or production with fixed costs, it is neither. No assumptions are made on the production function, and unlike fixed costs, which are usually treated as exogenous, a is endogenously chosen by the firm to affect worker utility and thus wages.

have a lower marginal cost, a sizable literature has found that larger firms pay more. This fact is not sufficient to reject the competitive model, as there are many alternative reasons for this relationship, like larger firms paying efficiency wages or targeting better workers, but the story is more complicated than this model in isolation. As we will see, much of the weight of the literature points to differences in firm productivity as the relevant omitted factor here.

3 Benchmark Comparative Statics and Other Sources of Amenity Variation

Strictly speaking, in the representative agent model there is no variation in amenity provision. However, comparative statics are informative about who receives amenities, and why they vary in practice. These comparative statics are based on the compensated demand for amenities, obtained from the expenditure minimization problem via Shephard's Lemma:

$$a_k^*ig(\overline{V}(s),c_1,\ldots,c_n;tig) \ = \ rac{\partial Cig(\overline{V}(s),c_1,\ldots,c_n;tig)}{\partial c_k}.$$

In this section I review how optimal amenities vary with amenity cost, preferences, and target market utility. I conclude with a discussion of other determinants in amenity variation that are outside of the representative agent model. In the next Section I discuss the Rosen model which gives a more complete account of amenity determination in an equilibrium model that allows for multiple dimensions of heterogeneity.

3.1 Variation in Cost of Amenity Provision

In the benchmark model, if we take as given that workers are homogeneous, and there is a reservation utility that has to be met, the source of variation in amenity provision is due to cost differences by firm. Compensated demand for amenity a_k is decreasing in c_k because the Hessian of $C(\cdot)$ is negative semi-definite. If we zoom out enough, this is an obvious source of actual variation in some dimensions of working conditions. The cost of a safe workplace is simply higher in construction than, say, the actuarial sciences. There is considerable evidence that costs affect the provision of amenities. For example, Oyer (2008) shows that firms are more likely to offer benefits that are related to their industry, presumably because they can do so at lower cost. By the same logic, larger firms are more likely to provide benefits. Figure 1 shows the optimal bundle

Figure 2: Public Good Amenity Determination (n = 1)



Panel A: Corner Solution

(point b) when the cost of the amenity declines to c' from c. The new isocost curve is the dashed line.

Tax treatment of benefits also plays an important role in amenity provision. Consider health benefits that a worker values according to a demand (MWTP) schedule D(a), where a is the benefit (D' < 0). The schedule is downward sloping because these benefits are not valued 1-1 with dollars, as they are restricted to certain uses. It costs the firm \$1 to provide \$1 of benefits. Absent regulatory or other market factors, workers are willing to trade benefits for cash up to D(a) = 1 - t. The prediction from this model is that the last dollar of benefits are valued by workers at the net-of-tax rate. This beneficial tax treatment motivation for non-wage amenity provision extends beyond traditional benefits. The value to a worker of any type of improved work conditions (e.g., investment in morale) is untaxed. By the first-order condition, a one dollar investment in working condition dimension k is net-beneficial for the firm so long as the worker valuation of that investment exceeds $(1-t)c_k$ dollars. The marginal value of a workplace investment is therefore determined by the tax rate. If tax rates increase, firms will spend more on all untaxed benefits as well as measures to improve working conditions. Arold et al. (2024) test and find support for this prediction in union contracts. In the context of Canada, they find that higher taxes are associated with more worker-rights clauses and lower pre-tax wages.

3.2 Variation in Preferences Over Amenities

Substantial preference heterogeneity over amenities, documented by Mas and Pallais (2017), among others, is another reason why amenity provision will vary. The benchmark model implies that if a worker places high value on the amenity, all else equal the firm will produce more of it. In Figure 1, this would correspond to a shift of the indifference curve along the firm's isocost line. One experiment to test this prediction causally would ask, if there is an exogenous change in the firm's workforce, does the mix of amenities it offers change? While I am not aware of an experiment like this one, Corradini, Lagos, and Sharma (2022) consider the opposite scenario where they find that firms that were effectively mandated to offer more "female-centric amenities" through collective bargaining agreements in Brazil had increases in the share of women in their workforce, due to both more female hires and fewer female separations.

3.3 Variation in Market Utility

A third reason amenities vary in this model is due to variation in reservation job value $\overline{V}(s)$. For homogeneous preferences, there is a utility level for each skill that clears

the labor market. These are given by $\overline{V}(s)$, which we assume is increasing in *s*. Firms hiring at higher skill levels will target higher utility levels according to this schedule. Figure 1 shows two bundles for $\overline{V}(s_1) < \overline{V}(s_2)$ with optimum bundles at points a and c. As the figure shows, if a_k are normal goods, as reservation value rises there is an overall increase in expenditure in both wages and non-wage amenities. This prediction is certainly borne out in the data as higher-earning workers also tend to have access to more benefits. This observation is relevant but often ignored when estimating Mincer models. Maestas et al. (2023) show that in their data there is a 0.52 log point difference in wages for workers with and without a college degree. Using their measure from a stated preference study on the value of a collection of working conditions (like scheduling flexibility and working from home), they estimate a 0.69 log point difference in total compensation.⁸ They do not consider fringe or other monetary benefits, so this divergence will be larger if we were to include them given existing evidence on differences in benefits by education (e.g., Pierce 2001).⁹

The observation that the share of total compensation in the form of benefits or other amenities is increasing with income is not implied by the representative agent model unless we make further assumptions on worker preferences or firm costs. For example, homothetic $V(\cdot)$ (like Cobb-Douglas) implies that the fraction of total expenditure on each amenity is constant as overall compensation goes up. An income elasticity of amenities that exceeds 1 can arise either from non-homothetic utility or if amenity costs decline in income. A salient example of the latter is the progressive income tax which progressively increases the price to the firm of providing after-tax wages to workers, and progressively lowers the relative price of providing non-cash amenities—which are often untaxed—to workers at higher tax brackets. Lamadon, Mogstad, and Setzler (2022) structurally estimate how a progressive income tax inefficiently increases employment in high amenity firms. (See Woodbury 1983 for earlier evidence.)

Variation in $\overline{V}(s)$ can also be used to model the effects of rising incomes. Hamermesh (1999) shows that night work fell sharply from the early 1970s to the early 1990s. He argues that this is the result of night work being a disamenity and an inferior good coupled with rising incomes. In the model this would be represented by rising amenity levels due to rising $\overline{V}(s)$.

While I have described the problem of a varying utility in terms of market-determined targets for skill or income, the conclusions follow for any variation in this target regardless of the reason. With market imperfections, the target job value will incorporate

⁸This estimate is a combination of workers with college degrees having more amenities and also placing a higher value on them.

⁹On the other hand, Mas and Pallais (2020) make a cautionary point that certain amenities are systematically bundled with disamenities, like stress on the job and workload.

turnover considerations or efficiency wages, and will vary across firms as in Sorkin (2018). This variation will generally translate in wages and amenities positively comoving which, in practice, confounds the negative relationship between wages and amenities (via varying amenity costs). There is a lot of evidence, that I will review, that variation in $\overline{V}(s)$, both due to skill and due to imperfect competition, is a dominant joint driver of wages and amenities.

3.4 Other Mechanisms

There are numerous mechanisms outside of the representative agent model that influence non-wage amenity provision. I will review a number of them here.

3.4.1 Productivity

On the spectrum of non-wage amenities are those whose purpose is just consumption value, like a cash benefit, and those that are primarily productivity-enhancing, but which may incidentally affect worker utility. In fact, just about every feature of the workplace can be thought of entering directly into a worker's utility function in some form. Investments in productivity, therefore, should be considered a driver of non-wage amenities. This can range from offering laundry services for employees, to encourage them to stay in the office, to investing in more efficient and quieter machinery. Oyer (2008) shows that long-hours are a strong predictor of whether an employer offers "effort packages"—benefits that are designed to lower the cost of higher effort at work, like meals and entertainment options at the workplace. Variables traditionally associated with higher benefits, notably wages, do not explain these effort packages. This finding indicates that effort packages are primarily used to motivate workers rather than for their consumption value.¹⁰

An implication of the close connection between productivity and amenities, is that capitalization of amenities into wages can potentially rationalize productivity differentials between firms, as surveyed in Syverson (2011). If a workplace amenity reduces productivity, firms can pay lower wages and remain competitive. For example, a retail store with many partially occupied employees can coexist next door to a store with fewer but highly occupied employees, while selling the same products at the same price if the inefficient staffing arrangement is itself a work amenity, via a more relaxed pace,

¹⁰Highlighting the fuzzy boundaries between productivity and cost initiatives and workplace amenities, Freeman and Kleiner (2000) discuss "Employee Involvement" programs. They note that these programs were touted as raising productivity and profits, but in their empirical analysis, they find negligible productivity effects but improvements in worker well-being.

thus allowing the store to pay lower wages. This story can extend beyond two firms to entire economies.

Owing to the importance of the question, there are hundreds of studies across disciplines (with varying degrees of rigor) on variants of the relationship between work conditions, amenities, or work environment on productivity. Themes that will be discussed in this Chapter include the literature on working from home (WFH) (e.g., Bloom et al. 2015; Emanuel, Harrington, and Pallais 2023) and workplace safety (e.g., Gray 1987). The representative agent model is silent on whether an amenity enters into the production function, but with perfectly competitive markets that requires the firm to stay on a target utility level it will not matter for how we analyze, for example, compensating differentials. All productivity benefits or costs will be reflected in output rather than the wage. If there are productivity-boosting amenities that workers enjoy without diminishing utility, the firm will just purchase as much of them as it can until they are at a corner solution.

3.4.2 Dynamic Considerations

Goldin and Katz (2011) and Goldin, Pekkala Kerr, and Olivetti (2020) highlight dynamic factors that infuence the provison of workplace flexibility policies, particularly for women. Women with higher educational attainment who plan to have kids will value these benefits more than those with lower educational attainment because they have a higher monetary return to work, and these policies make it easier for them to continue their careers. At the same time, theoretically, firms with a higher education workforce are more likely to offer these policies because the cost of offering, say, paid leave is easier to recoup since women with more education are more likely to return to work.

The dynamic problem of the workplace flexibility case is more general. Amenities are often investments. They require possibly large fixed costs and are recouped over time through whatever are the relevant benefits, such as lower wages, less turnover, fewer vacancies, or higher productivity. While I consider the static case, it is easy to see how dynamic considerations lead to additional predictions on the determinants of amenity provision if we think of amenities in this way. Interest rates, credit availability, policy stability, and growth will all be relevant factors in determining changes to working conditions.

3.4.3 Unions

The theory of fringe benefit determination in unionized workplaces is discussed in Freeman (1981). The link between unionization and the provision of fringe benefits is clearest for craft unions. Because their members, like carpenters and plumbers, are attached to occupations rather than firms, unions provide the institutional capacity to operate fringe benefit programs.¹¹ Trade unions may increase fringe benefits because they negotiate higher total compensation, and benefits are income elastic, or because the union's preference is determined by its median voter, who may have a higher preference for benefits than the marginal worker who determines benefit provision in non-union firms. Freeman empirically finds that unions have sizable impacts on the provision of fringe benefits. In the context of Brazil, Lagos (2024) finds that an increase in union bargaining power, via a judicial order, led to significant impacts on amenities in collective bargaining agreements. Approximately half of the overall gain in the total value of the collective bargaining agreements was through higher amenities.

3.4.4 Inside the Firm Heterogeneity and Spillovers

While workers imperfectly sort across firms in terms of skill and preferences over amenities, benefits are often provided uniformly in firms. A more realistic way to model benefits is to allow for different preferences over amenities, or different amenity demands, in the same firm. Inside the firm, heterogeneity is relevant for analyzing the provision of benefits, the distribution of surplus, and the consequences of shocks.

There is no role for within-firm heterogeneity in the representative agent model. If the amenities are public goods, and sorting is imperfect, the firm has to determine an amenity mix accounting for a distribution of preferences. As previously discussed, in the union case, benefit determination could be determined through a voting process such that benefit provision reflects the median union member. In the non-union case, the employer can maximize utility subject to an overall expenditure, over the entire workforce or some segment of the workforce. Since benefits tend to be income elastic, a reasonable hypothesis is that they are oriented towards top-earners.

In principle, the firm can personalize private good amenities on some observable dimension like income. However, there are constraints. Within-firm provision of benefits is best understood in light of legislation that seeks to make health and retirement benefits broadly accessible in workforces and prevent their concentration among top earners. In the context of retirement benefits, the Employee Retirement Income Security Act (ERISA) enforces nondiscrimination rules to ensure employer-sponsored plans, like 401(k)s and pensions, provide equitable benefits across income levels. These rules prevent plans from disproportionately favoring highly compensated employees (HCEs) over non-highly compensated employees (NHCEs) through annual testing, such as the

¹¹From this perspective, programs like the ACA potentially weakened these unions' advantage in benefit provision.

Actual Deferral Percentage and Actual Contribution Percentage tests. To comply, employers may need to reduce contributions for HCEs if NHCE participation or contributions are too low or provide additional contributions to NHCEs (called "qualified non-elective contributions"). The implication is that these rules compress benefits. The presence of HCEs and NCHEs together can make benefits more generous for low earners and less generous for high earners, relative to the case in which a firm employs only high or low earners.

Health insurance benefits are another example of a one-size-fits-all approach at most employers, due both to (likely) fairness considerations and nondiscrimination terms in IRS code for the tax deductibility of benefits that prevent employers from favoring highly compensated employees.¹² Like retirement benefits, a reasonable hypothesis is that health benefits are optimized to recruit high-income workers, but are then made available to lower-income workers as well. The cost of these benefits, then, are a floor for the marginal cost of hiring low-income workers. In theory, wages could be reduced to equalize total compensation to the value without benefits, but this requires that low-income workers have the same valuation for the benefits. Given tax progressivity considerations and diminishing marginal utility of income, this may not be the case. If these benefits are not valued at cost, the employer faces a wedge that is equivalent to a head-tax: they have to pay some premium over every dollar of compensation to a low-income employee.¹³ There is a perspective that this aspect of employer-sponsored health insurance has been devastating for the lower part of the income distribution (see, e.g., Case and Deaton 2020). Finkelstein et al. (2023) employ an incidence analysis and conclude that the college wage premium would be substantially reduced if the current employer provision of benefits were replaced with a payroll tax.

Spillovers of benefit provision within firms are empirically relevant. Freeman (1981) reports sizable spillovers from union to non-union workers in firms. Ouimet and Tate (2023) is one of few studies that has benefit information linked to matched employeremployee data. Using detailed job-level data from the BLS' National Compensation Survey matched to the LEHD, they document that nondiscrimination rules result in significant compression of retirement and health benefits within firms. They find that the

¹²Specifically, Internal Revenue Code Section 125 covers cafeteria plans, and Section 105(h) covers self-funded plans.

¹³While employers cannot offer less generous benefits to lower earners, they sometimes do offer lower premiums in lower pay bands. This is usually done to remain in compliance with the ACA, which states that for large employer's coverage to be deemed affordable it cannot exceed 9.5% of a household's income. When structuring the premium schedule this way, high earners are subsidizing low earners. From the employer's perspective, this further increases the wedge between compensation and marginal cost because each low-earner hired has some impact on premium reductions to stay in compliance with the ACA.

presence of highly compensated employees drives up benefit levels across the firm, reflecting the regulatory and administrative tendency to equalize benefits across workers. For multi-establishment firms, they use an IV of the average health benefit of establishments in the same MSA and industry as a firm's headquarters for the average benefit in non-headquarter establishments. Using this instrument, they estimate the causal effect of health benefits on separations. They find that health benefits positively influence entry of high-income workers, and positively impact exits of low-income workers. The interpretation is that benefits are optimized to recruit at the upper end of the workforce distribution and reduce job surplus at the lower end. In this way uniformity creates inefficiencies for firms employing lower-wage workers who may not fully value these benefits. Consequently, firms may respond by reducing their reliance on low-skill labor, turning to automation or outsourcing. Because of the wedge between between benefit valuation and cost, joint surplus will be low for lower-income workers making them more susceptible to separations following shocks.

4 Two-Sided Heterogeneity and Sorting: The Rosen Model

While the representative agent model goes a long way in delivering the predictions that have been the focus of much empirical work on amenities, and can handle comparative statics when there is variation in either costs or preferences, it cannot speak to the case of heterogeneity in both preferences and costs. The methodological hallmark of the Rosen model is characterizing equilibrium when there is two-sided heterogeneity in the market and sorting.

The formal origins of modern models trace to Lewis (1969), Rosen (1986) and Rosen (1974). Lewis (1969) developed a model of wage determination when workers and firms have preferences over work hours, which he treated as an amenity. Rosen formalized and expanded this model into what is now the canonical model of hedonics.

In the most general form of the model there is labor, a numeraire, and a continuum of firms that can be decomposed into their attributes $a = \{a_1, \ldots, a_n\}$. Each of these attributes (synonymous with amenities) is valued by workers and are costly for firms to produce. The production technology of a firm allows it freely choose to offer either nothing or a single job. But it needs to use as input a unit of labor plus a cost that depends on the amenities chosen. Workers are endowed with one unit of labor and have preferences over the jobs and different amounts of the numeraire. In equilibrium, the hedonic wage function has the form $w^*(a_1, \ldots, a_n)$. By definition, when the firms profit-maximize and the workers utility-maximize given $w^*(a_1, \ldots, a_n)$, the markets clear. That is, there may be inactive firms, workers that got out of the market, and matched pairs that will satisfy the self-selection property: they would not want to switch to the deal offered by another firm or worker. This is a Walrasian equilibrium, and hence Pareto optimal. In equilibrium there is perfect sorting. All workers choosing a specific bundle are identical. One of the main messages of Rosen (1974) is that the equilibrium hedonic wage function $w^*(a_1, \ldots, a_n)$ does not have an easy interpretation in terms of the primitives of worker or firms.

In what follows, I will outline the Rosen model so that it is easier to understand the logic of the equilibrium. I start with a simplified version that expands the representative agent model I presented to have multiple worker and firm types. I then outline the discrete version of the model presented in Rosen (1986) as it is frequently used in applications. Because of the complexity of the model, I will work through cases of a single amenity. Because of the multi-margin problem discussed above, this single amenity is best thought of as a composite index.

4.1 A Worked Out Case with a Continuous Amenity

To move from the representative agent to the Rosen model we have to add heterogeneity in both workers and costs. In doing so, we have to maintain an equilibrium whereby the firms of all types are profit-maximizing, and no worker wants to change to a different (a,w) bundle that is offered in the market. As in the representative model, I will consider a continuous amenity a, but introduce a large number of discrete firm and worker types.

Worker utility is defined as $V(a, w; \alpha_i)$, for $i \in \{1, ..., I\}$. The utility function determines the strength of preference for the amenity. The *i* subscript determines a worker preference type, and higher numbered subscripts have a stronger preference for the amenity. Firms cannot observe worker type, can offer only one (a, w) bundle and, for the time being, can only hire one worker. Firm profits are given by $\pi = p - w - c_j a$, where the *j* subscripts denote firm cost types. The marginal cost of amenity provision c_j is decreasing in *j* (higher types make it more cheaply). The market is perfectly competitive and firms are price-takers. Labor is supplied elastically at each preference type.

I will describe graphically an equilibrium that is incentive-compatible and results in perfect sorting of types. All workers and firms types are ranked, and in equilibrium each worker type matches to the firm type of the same rank. The solution is iterative. We start with type-1 firms and workers, then solve for type-2, and we conclude with type-3. The solution for all higher numbered types will follow from the same logic.

In equilibrium type-1 workers (those with α_1 preferences) match to firms in a no-

amenity sector (those with c_1 costs) and receive wage w_c . The wage w_c is determined competitively. This is point $(0, w_c)$ in Panel A of Figure 3.

Next we determine the allocation for type-2 workers. The type-2 problem is simply the representative agent benchmark where type-2 workers match to type-2 firms with a market utility given by the utility type-2 workers obtain in the no-amenity sector. The determination of their (a, w) bundle is shown in Panel A of Figure 3. Note that the relevant type-2 indifference curve is the one that passes through w_c , which corresponds to $V = V(0, w_c; \alpha_2)$, such that they have no incentive to deviate. The solution is denoted point a in the graph, with wage w_2 . Since the type-2 firm isocost intercept is below w_c the marginal cost to hire a worker for type-2 firms is lower than in the no-amenity sector by the difference in the intercepts. Therefore, type-2 firms have strictly higher profits than firms in the no-amenity sector even though they have the same production technology and the amenity is non-productive. These profits are Ricardian rents that arise from firms that are endowed with lower costs to provide the amenity, which allow these firms to compensate workers at an equal utility level as firms without amenities, but at a lower average cost.

Panel B of Figure 3 shows the determination of the equilibrium bundle for type-3 workers and firms. We first plot the contract curve between these type-3 workers and firms, along which there are no gains from trade. Next we ask, what is the profitmaximizing point on the contract curve such that no types deviate? These conditions are met at point c, where the contract curve intersects with the type-3 indifference curve that is tangent to the profit-maximizing type-2 isocost curve. Neither type-2 workers nor firms will do better by moving to this bundle since it is above the profit-maximizing type-2 isocost and below the type-2 indifference curve. If type-3 firms offer any bundle that is lower on the type-3 contract curve, no workers will join them because a type-2 firm can offer the bundle at point b and hire type-3 workers at the same cost as type-2 workers.¹⁴ Equilibrium bundles for all types greater than 3 follow by the same logic. Equilibrium type-k bundles are on the type-k contract curve, at the intersection of the the type-k indifference curve that is tangent to the type-(k-1) isocost curve. An implication of the iterative form of this equilibrium is that optimal bundles for type k are a function only of the amenity costs and worker preferences types less than k.¹⁵ A regulation affecting high-cost (low-type) firms will have a ripple effect across the entire labor market. Whereas the effects of a regulation on low cost (high type) firms will be more

¹⁴In this scenario type-2 firms will be indifferent between hiring at points a and b and will choose between these bundles using a tie-breaker rule so as to not violate the one offer per firm condition.

¹⁵While this simplified version of the Rosen model reads like a screening model, information asymmetries do not lead to inefficiency because the identify of who matches is irrelevant. If firms observe their worker types, nothing changes. The market is competitive and the first welfare theorem holds.

Figure 3: Amenity Determination with Two-Sided Heterogeneity and Sorting





insulated.

Because the type-3 isocost has a lower intercept than the type-2 isocost, we know that they have a lower cost per-worker and higher surplus relative to type-2 firms. If we introduce an employment margin and firms equate their marginal revenue product of labor to their expenditure per worker, a lower price per worker translates to higher employment. Thus, low-amenity cost firms are both more profitable and larger.

The surplus for the amenity endowment is shared between workers and firms. For k > 2, Type k workers have surplus if we define it as a positive utility difference relative to their next best alternative. Because the type-3 bundle is defined based on a tangency with the type-2 isocost curve, type-3 workers would do strictly worse by moving to the type-2 bundle. There is also surplus if we define it as the difference in wages relative to the no-amenity sector $(0, w_c)$, which is what the type-3 workers "paid" for the full amount of the amenity, versus the WTP for the amenity. We know this because in equilibrium $V(a_3, w_3; \alpha_3) > V(0, w_c; \alpha_3)$.

High amenity cost (low-type) firms are at an inherent disadvantage in this market because they have the same production technology as other firms but have to pay a higher cost to recruit workers. They are not competed out of the market because they can offer high wage and low amenity bundles that are attractive to some workers who value more cash, and sub-optimal for higher type firms to offer. The reason, then, that firms with different costs co-exist in this model is because a firm can offer only one compensation bundle. Otherwise, a type-3 firm can mimic a type-2 firm and offer bundles that are marginally more attractive to type-2 workers at a lower overall cost than the type-2 firm bundle. The fact that we do observe heterogeneity in amenity provision in reality suggests that firms are constrained in their ability to offer different amenity levels. Outsourcing might be seen as a way for an amenity-rich firm to offer high-wage and low-amenity positions that its high-amenity preference workers will not accept. An example might be the high utilization of travel nurses, who as contract employees at hospitals tend to earn higher wages than staff nurses but do not have the same employment protections or fringe benefits (see Seo and Spetz 2013).

4.1.1 Long-run Equilibrium

The Rosen model I described is in the short-run. It asks what is the equilibrium outcome given an existing set of firms with a distribution of amenity costs. A long-run model with no barrier to entry leads to the degenerate case that, in a given market, only the set of firms with the lowest cost to produce the amenity survive. To see this, we can impose a zero-profit condition. The hedonic wage curve in the zero-profit case is entirely determined on the production-side by the upper-envelope of the zero-profit isocost curves. Rents are then fully transferred to workers, but since rents are higher in higher amenity (lower amenity cost) firms, this is not an equilibrium. Workers in high amenity cost firms will want to move. The long-run equilibrium then is for all workers to move to the lowest amenity cost firms, thus leading to the degenerate outcome. Put more simply: firms with a low cost to produce amenities have a competitive advantage because they can make something workers like more cheaply. If their profits are driven to zero in the long-run, firms with higher cost structures must face closure. With only one firm type remaining, workers can still vary in wage/amenity bundles as they position themselves along the isocost curves for these remaining firms. Note this long-run equilibrium holds only within a labor market, defined for example by industry. In the long-run there will still be differences in costs across industries.

The competitive long-run prediction of no variation in amenity costs may have some empirical relevance in that the basic benefits package of health insurance, 401(k), and paid leave, all of which are easy to replicate, are available in most jobs for collegeeducated workers. Moreover, the share of employment at large firms has been steadily increasing according to Census Bureau's Business Dynamics Statistics data, consistent with the long-run prediction as larger firms tend to have more amenities and better working conditions.

However, it is obviously untrue in general that workplace conditions are uniform within a market. How, then, do low amenity workplaces survive when there is free entry and worker mobility? One possibility is that the amenity cost is a function of an inelastic factor. For example, there might be a fixed set of managers that can implement safe and friendly workplaces at relatively low cost. In this scenario, in the long run, the amenity variation between firms will remain because not all talented managers can match to a firm.¹⁶ A zero-profit condition will mean that the wage of good managers will be bid up. In this variation of the model, rents from better amenities are captured by these managers.

Another explanation that allows for equilibrium variation in amenity provision and free entry of firms in the long-run is any barrier to entry on the worker side. For example, unions are uniquely suited to ensure certain valuable amenities for workers, such as grievance procedures, collective advocacy, and worker voice more broadly. Unions lower the cost of amenity provision to firms, and restrict worker entry to ensure that their members capture those rents. Another example for how amenity variation can be

¹⁶An executive at a large supermarket chain told me that the reason why injuries had fallen so quickly at their company in the previous five years was because "it was only after we decided it was a big enough problem that we assigned our best people to solve it." This response highlights the management talent bottleneck.

sustained in the long-run is spatial. While there is free entry of firms into sunny California, housing restrictions curb labor supply. In this case, homeowners, rather than workers, capture all of the amenity rents via these barriers.

4.2 Discrete Case

The discrete amenity version of the model is developed by Rosen (1986). It has persisted in the labor literature because it allows for transparent analysis of a highly complicated environment with different dimensions of heterogeneity and sorting and generally gives the right intuition, particularly on the importance of the price depending on the marginal worker and firm.

In the model there is a discrete amenity, such as a dangerous (d = 1) or safe (d = 0) workplace. The elements of this model are workers (indexed by *i*) defined by their heterogeneous willingness-to-accept (WTA_i) to join dangerous workplaces, and firms (indexed by *j*) who are defined by their heterogeneous marginal cost to invest in safe workplaces (MC_j). We define these distributions of heterogeneous worker preferences and firm costs as WTA_i ~ $G(\cdot)$ and $MC_j ~ F(\cdot)$, respectively. These primitives are sufficient to obtain a market compensating differential between jobs with and without the disamenity.

Define the difference in the wage between dangerous and safe jobs as $\Delta w \equiv w_{d=1} - w_{d=0}$. Workers accept a dangerous job if WTA_i < Δw (extra pay exceeds requirement) and firms maintain a dangerous workplace if MC_j > Δw (wage savings are lower than cost of safe workplace). In this model, the compensating wage differential is at the intersection of the supply and demand functions for the amenity, which is shown in Figure 4. The inverse supply curve of workers into dangerous jobs plots $\Delta w = G^{-1}(D_w)$, where D_w is the fraction of workers in dangerous jobs. When dangerous jobs pay more (higher Δw), more workers find it advantageous to work in dangerous jobs. The inverse demand curve of firms for dangerous jobs plots $\Delta w = F^{-1}(1 - D_f)$, where D_f is the fraction of firms with dangerous jobs. When dangerous jobs have a higher premium (higher Δw), more firms pay for safety and there are fewer dangerous jobs. The market clears when the wage differential equates these two shares.

While simple, the model has several interesting implications. The first is that the market compensating differential for a job with an amenity relative to the job without the amenity is pinned down by the MWTP of consuming the amenity by the worker and marginal cost of producing the amenity by the firm. At the equilibrating point, these are the same. This implication is not innocuous because the marginal worker and firm may differ substantially from the average, something that needs to be considered



Figure 4: Supply and Demand for Discrete Amenity

Note: The supply curve traces worker supply to dangerous jobs. The demand curve traces firm demand for danger dangerous jobs.

when using these measures for evaluating regulation. For example, over the last two decades workplaces have become much safer. A plausible explanation for why injuries have fallen is that firms invested in safer workplaces in response to rising health care and workers compensation costs, which made injuries more expensive to firms and increased the net benefit of making safety investments. The Rosen model predicts in this case that the compensating differential for a workplace with safety improvements should fall (the demand curve in Figure 4 shifts left), not because workers valuation of safety has changed, but because the marginal worker in a dangerous job now has a higher risk tolerance. If all we observed is the wage differential, we might draw the wrong conclusion.

A second observation from this discrete amenity model is that inframarginal firms and workers have surplus in the sense that for a small change in the wage their decision will not change. Inframarginal workers with a disamenity will have higher utility than marginal workers because they are paid the same, but have more tolerance for the disamenity. Inframarginal firms with an amenity have higher profits than marginal firms because their cost to providing the amenity is lower. Worker surplus in this discrete amenity case comes about because of imperfect sorting. There is only one wage, but workers choosing the same amenity have heterogeneous preferences, so some have to be inframarginal.

4.3 Evidence on Sorting

The Rosen model makes a stark prediction that there is sorting between workers and firms on their respective preferences and costs. How well does this prediction hold? The answer to this question is relevant both to test the theory of equalizing differences and to determine the potential surplus workers have in a job, defining surplus in this case as the gap in utility between a job and their next best alternative. When there is perfect sorting there is limited room for amenities to generate surplus. But with imperfect sorting, and if amenities are determined at the firm level, inframarginal workers who have the highest valuation over amenities will have surplus.

I characterize the evidence on sorting as on balance supportive of the prediction from the equalizing differences model, but it depends on the domain considered, and is certainly far from perfect. On the plus side, Krueger and Schkade (2008) find a strong relationship between how social someone is outside of work and whether they are in a job that involves interacting with other people. Using a model of occupational choice, DeLeire and Levy (2004) find that single parents have elevated aversions to risky occupations relative to other workers of the same gender. Their bottom-line is that differences in occupational risk explain one-quarter of occupational gender segregation. In the context of military occupations, Greenberg et al. (2021) find that the indifference curves relating wages to mortality risk for women and men in noncombat roles are substantially steeper than those for men in combat occupations, indicating a stronger aversion to mortality risk in these subgroups. There is evidence that people who have a higher WTP for more family-friendly work arrangements are more likely to be in jobs with these arrangements (Mas and Pallais 2017; Wiswall and Zafar 2018). In Mas and Pallais (2017), workers in formal WFH arrangements have approximately twice the WTP for WFH than workers not in these arrangements. Moreover, women with children under 4 have both elevated WTP for formal arrangements and actual rates in these jobs. As previously discussed, Corradini, Lagos, and Sharma (2022) show that when there is exogenous implementation of "female-friendly" amenities, the share of females employed in the firms increases.

On the negative side, Lachowska et al. (2023b) find that there is very limited sorting between workers who prefer long hours and firms that offer long hours. Their explanation is rooted in imperfect competition: workers in general experience upward equilibrium hours constraints (they want more hours than they have), firms that offer overall better compensation do so both on the wage and hours margin, and these more attractive firms select workers from queues on dimensions other than preferences. In the context of work arrangements, while Mas and Pallais (2017) document that people who have a high WTP to avoid irregular, employer-driven schedules are not less likely to be in those jobs.

5 Hedonic Models: Estimation

The requirements for estimating hedonic models depends their intended purpose. They are less stringent if the goal is to estimate an equilibrium relationship, more stringent if the goal is to obtain MWTP estimates, and even more stringent if we want to estimate indifference curves.¹⁷

For any purpose, in the best world we have an amenity cost shifter that changes the relative price of an amenity and then use that to instrument for the quantity of the amenity to estimate its effect on the wage. This estimate is interpretable as the derivative

¹⁷Rosen (1974) proposed a two-step procedure for estimating indifference curves–also called bid curves–using a hedonic model. In the first step a hedonic model is estimated to predict the price as a function of product attributes. In the second step the derivative of the estimated price function is regressed on a non-linear function of the the amenity of interest to recover the indifference curve. This procedure has not been widely adopted due to the very strong parametric assumptions required for identification (Greenstone 2017).

of the hedonic wage function w(a) at some initial amenity level. The cost shifter moves the firm's isoprofit along this equilibrium schedule as it trades-off the amenity for the wage. As I discussed earlier, if we had such a shifter, the first-stage of this IV regression is also highly informative for how firms and workers value incremental changes in the amenity since it measures movement on the amenity demand curve.

Owing to the rarity of experimental variation, much of the empirical literature has focused on what we can learn from analysis of observational data. While this approach has many potential pitfalls, there is much to be learned if we think carefully about the data generating process in equilibrium models. This is especially true for long-run steady-state relationships, about which experiments can typically tell us little.

5.1 Hedonic Models: Assumptions

Given observational data on amenities and wages, it is helpful to ask what is the data generating process, and particularly the sources of heterogeneity that produce these values. In the Rosen model, within a market, this variation comes from heterogeneity in worker preferences and firms' costs. Any correlation of wages with amenities will reflect preference and cost parameters and will not be particularly informative about the primitives. We can appreciate this problem in Panel B of Figure 3, where the observed hedonic equilibrium are the equilibrium bundles $(0, w_c)$, (a_2, w_2) , and (a_3, w_3) . Even if one could hone in on a specific market, with appropriate data, and adequately control for unobserved skill, the slope of this equilibrium relationship does not identify worker or cost parameters because over any range of amenity values the slope of the hedonic curve mixes the slopes of the indifference curve, the isocost curves, and changes in composition of workers and firms.

A healthy perspective, embraced by Lavetti and Schmutte (2018), is that we should be satisfied to obtain an unbiased estimate of the market hedonic price function. This function is itself an interesting equilibrium object which tells us how amenities tradeoff in the market equilibrium, even if the slope of this function is not interpretable in terms of worker or firm primitives. The hedonic price function is the relevant object, for example, if we are interested in understanding how variation in amenities affect the overall wage structure, a question that has recently been in focus (Sorkin 2018; Taber and Vejlin 2020; Lamadon, Mogstad, and Setzler 2022; Morchio and Moser 2024).

For purposes of discussing assumptions in estimation it is helpful to return to the more tractable representative agent model. This model gives an expression for the equi-
librium wage of the form

$$w^*\big(\bar{V}(s),\mathbf{c}\big) = C\big(\bar{V}(s),\mathbf{c}\big) - \sum_{k=1}^n c_k a_k^*\big(\bar{V}(s),\mathbf{c}\big) + \varepsilon,$$
(4)

where $C(\bar{u}, \mathbf{c})$ is the expenditure function, and \mathbf{c} is the vector of amenity costs. I have added an i.i.d. error term ε to reflect that wages are now an object we observe in data. This expression just says that the wage is the minimum expenditure the firm incurs given amenity costs to match reservation utility, less the optimal expenditures on amenities. The wage depends on the firm's job value reservation target, and the full set of amenity costs. I detail some points on estimation below:

- **Unobserved skill** The expenditure term $C(\bar{V}(s), \mathbf{c})$ is typically unobserved. It is increasing is skill *s* because higher-skilled workers have more total compensation. Skill affects amenity provision for a fixed amenity cost-it is an argument in $a^*(\cdot)$ —but it can also be correlated to amenity cost through sorting of workers to firms. Any of these factors will result in an amenity term that is correlated to the error, resulting in biased estimates of the slope coefficients on the amenities. Much attention has been given to this issue; see, e.g., Brown (1980) for an early discussion and Kostiuck (1990), who finds evidence of strong self-selection of workers with low potential earnings into night work. A novel approach for addressing unobserved individual heterogeneity is Stern (2004), who looks at newly minted biology PhDs who receive multiple, simultaneous job offers. By comparing these offers within each candidate, the study controls for unobserved heterogeneity in ability and preferences. Within candidate, there is a negative relationship between science-oriented jobs and compensation, implying that the biologists are paying for research autonomy via lower salaries. This negative relationship does not hold without researcher fixed effects, implying that employers vary compensation in their offers, and candidates who tend to command higher wages regardless of job type are more likely to receive science-oriented offers.
- **Unobserved attributes** Workplaces are complex and multi-dimensional. Equation (4) assumes we can observe all attributes, but this is implausible. In practice, good measurement of non-wage amenities beyond traditional benefits presents a major challenge to implementing these types of models and much of the frontier research is on finding new ways to estimate workplace amenities and working conditions. Notable example using textual analysis are Arold et al. (2024) and Lagos (2024) on collective bargaining agreements, and Sockin (2022) on employer review sites. Absent comprehensive measurement, experimental variation

in an amenity or its cost would address this problem if the goal is to isolate the relationship of a single amenity to the wage.

- **Firm pay policy** The expenditure term $C(\bar{V}, \mathbf{c})$ will also capture any variation in firm pay policy that are predicted by most models of imperfect competition, for example rent-sharing, efficiency wages, or collective bargaining. Any variation in $C(\bar{V}, \mathbf{c})$ across firms at a fixed level of skill will result in amenity values that that are correlated with the error both because target utility is an argument in $a^*(\cdot)$ and because there can be an equilibrium relationship between reservation utility and amenity cost. In the Hwang, Mortensen, and Reed (1998) equilibrium, firms with lower endowed amenity costs have higher target utility because there is a higher opportunity cost of having vacancy. The bulk of the evidence suggests that this is an important channel. To my knowledge, it is uniformly the case across studies that measures of firm- or establishment-level pay are positively correlated with firm amenities, and moves up the job ladder increase wages and amenities in tandem (Lavetti and Schmutte 2018; Lachowska et al. 2023b; Sorkin 2018; Roussille and Scuderi 2023; Sockin 2022; Lamadon, Mogstad, and Setzler 2022, to name a few.).
- **Sticky wages** If we have experimental variation in amenity provision or costs, to interpret the wage response as a MWTP we have to assume strict adherence to competitive reservation utility. However, wages are sticky (Quach 2025). It is more likely that we observe a hedonic relationship over the medium- to long-run rather than immediately after a shock. By implication, when estimating compensating differentials using panel data, using job changes is preferable to using variation in amenities over a match.
- **Multiple-Margin Adjustment** I have already discussed this problem, but it is worth repeating. The assumption that all other amenities are held constant when when amenity changes is incompatible with firm optimization, generally. This severely limits one's ability to read off a compensating differential, even if it experimentally generated, as a MWTP. One way to proceed is to accept that wage pass-throughs are interesting on their own even if they do not necessarily have a structural interpretation. If we insist on the structural interpretation, then we can take the estimates as lower bounds, assuming no other biases. Another approach is to consider, when possible, the relationship between wages and bundles of attributes. For example, Tsao (2025) compares the valuation of a wide-class of amenities and working conditions (via stated preference surveys) for teachers and

their next-best outside option. She finds that the outside option tends to pay less, but it has higher overall amenity values. In this case, multi-margin optimization is less relevant because it is happening within the broader bundle being compared.

5.2 Compensating Differentials in Two Settings

Given the identification challenges described above, it is not surprising that so many of the published and unpublished estimates of compensating differentials from regressing wages on job attributes often have counterintuitive signs, fail robustness checks, or have implausible magnitudes.¹⁸ In this section I will talk about a two domains that are representative of the mechanisms at play. My objective here is not to provide a comprehensive survey of all work done to estimate compensating differentials, but to provide a somewhat high-level overview of how the empirical literature accords with theoretical predictions. I will begin discussing the provision of work from home (WFH) options and then will review the evidence on workplace safety. To foreshadow my conclusions, it is difficult to reconcile the available evidence without allowing for both selection on individual unobservables and deviations from the perfectly competitive benchmark, such as rent-sharing.

5.2.1 Working From Home as a Job Attribute

Much of the focus of the more recent literature on job amenities has been on alternative work arrangements to traditional 9-5 on-site work. Even before the COVID-19 pandemic, WFH was a focal attribute. The literature has estimated how workers value WFH, its productivity impacts, and its implications on spatial sorting. My high-level read on the conclusions of the literature on how WFH jobs are valued is:

- In cross-sectional hedonic regressions with many controls for observable characteristics including education, occupation and industry, WFH jobs are associated with a positive pay premium (Mas and Pallais 2017). This is true both before and after the COVID-19 pandemic, which led to a sharp rise in these positions. If we interpret this premium as a MWTP, it suggests that WFH positions are a disamenity.
- 2. However, there is considerable evidence from choice experiments and stated preference surveys that workers on average like WFH jobs and are willing to take a pay cut to be in one (Mas and Pallais 2017; Maestas et al. 2023; Aksoy et al. 2022).

¹⁸Lavetti (2023) discusses some of this evidence.

3. There is evidence that workers negatively select into WFH jobs based on their productivity (Emanuel and Harrington 2024).

What do we conclude from these three findings? Facts 1 and 2 in isolation suggest that there is a skill component that is unmeasured by the econometrician, but rewarded by the market, that is positively correlated to WFH. Unobserved skill results in positive bias of the WFH job attribute, thus undoing the compensating differential. Yet Fact 3 suggests that the opposite is happening. When given the option, workers negatively select into WFH arrangements based on individual productivity.

Collectively, these facts can be reconciled once we allow for job characteristics that are unobserved in typical datasets used to generate Fact 1. This job characteristic is both positively correlated to the wage (via higher productivity) and the probability of a WFH option. For example, such a characteristic might be job autonomy: whether the job gives discretion on how a worker performs tasks. By excluding an indicator for autonomy from the wage equation, its positive relationship with the wage loads onto the WFH variable, resulting in a positive bias.

Pursuing this mechanism, wage determination for worker i and firm j can be written as:

$$\ln(w_{i,j}) = \beta_w \operatorname{WFH}_{i,j} + \alpha_i + \beta_T \ln(T_{j(i)}) + \varepsilon_{i,j}$$

where $T_{j(i)}$ is a firm- or job-level productivity measure in firm *j* where worker *i* is employed. Person-specific wage heterogeneity, including unmeasured skill, is captured by α_i , and WFH_{*i*,*j*} is an indicator of whether person *i* in job *j* is working from home. If the firm is a strict price-taker, then $\beta_T = 0$ because firm-level productivity shifters will not affect wages it pays for each skill type in the labor market. Therefore, the interpretation of β_T is as a term that governs rent-sharing. Thus, the $\beta_T \ln(T_{j(i)})$ term can be interpreted as an AKM firm-effect.

If we estimate a short regression, omitting both α_i and $\ln(T_{j(i)})$:

$$\ln(w_{i,j}) = \beta_w^s \operatorname{WFH}_{i,j} + u_{i,j},$$

then by the omitted-variable bias (OVB) formula,

$$\beta_w^s = \beta_w + \frac{\operatorname{Cov}[WFH, \alpha_i]}{\operatorname{Var}[WFH]} + \beta_T \frac{\operatorname{Cov}[WFH, \ln(T_j)]}{\operatorname{Var}[WFH]}.$$

Assume

$$\alpha_i = \gamma_{\alpha, \mathrm{WFH}|T} \, \mathrm{WFH}_{i,j} + \lambda \, \ln(T_{j(i)}) + u_i$$

where u_i is uncorrelated with both WFH_{*i*,*j*} and ln($T_{j(i)}$). Then

$$\operatorname{Cov}[\operatorname{WFH}, \alpha_i] = \operatorname{Cov}\left[\operatorname{WFH}, \gamma_{\alpha, \operatorname{WFH}|T} \operatorname{WFH} + \lambda \ln(T_{j(i)})\right]$$
$$= \gamma_{\alpha, \operatorname{WFH}|T} \operatorname{Var}(\operatorname{WFH}) + \lambda \operatorname{Cov}[\operatorname{WFH}, \ln(T_j)].$$

Hence,

$$\frac{\text{Cov}[\text{WFH}, \alpha_i]}{\text{Var}[\text{WFH}]} = \gamma_{\alpha, \text{WFH}|T} + \lambda \frac{\text{Cov}[\text{WFH}, \ln(T_{j(i)})]}{\text{Var}[\text{WFH}]}$$

Substituting back into the short-regression formula:

$$\beta_{w}^{s} = \beta_{w} + \gamma_{\alpha, \text{WFH}|T} + \left(\lambda + \beta_{T}\right) \frac{\text{Cov}[\text{WFH}, \ln(T_{j(i)})]}{\text{Var}[\text{WFH}]}.$$
(5)

This expression says that the bias on the WFH estimate is a function of negative selection into WFH ($\gamma_{\alpha,WFH|T} < 0$), the relationship between unobserved skill and firm productivity (λ), and the relationship between the AKM firm effect in the log wage and WFH ($\frac{\text{Cov}[WFH, \beta_T \ln(T_{j(i)})]}{\text{Var}[WFH]}$).

We can get some additional insight using known information on these terms. There is an extensive literature on rent-sharing elasticities, and an estimate around $\beta_T = 0.1$ is reasonable (Abowd and Lemieux 1993; Card et al. 2018). Mas and Pallais (2017) estimate an average WTP for WFH of 8 percent, implying that $\beta_w = -0.08$ in the long regression where the coefficient on the WFH indicator represents pure compensating differentials.¹⁹ Mas and Pallais (2017) also estimate $\hat{\beta}_w^s = 0.1$. The parameter λ represents the slope coefficient of the regression of the individual component of log wages on log productivity. This value can be found from estimates of AKM models.²⁰ In Table 3 of Lachowska et al. (2023a), it is estimated as $\lambda = 0.86$. Given these values, we obtain:

$$0.1 = -0.08 + \gamma_{\alpha, \text{WFH}|T} + (0.86 + 0.1) \frac{\text{Cov}[\text{WFH}, \ln(T_{j(i)})]}{\text{Var}[\text{WFH}]}.$$

Because $\gamma_{\alpha,\text{WFH}|T} < 0$, it follows that

$$\frac{\operatorname{Cov}[\operatorname{WFH}, \ln(T_{j(i)})]}{\operatorname{Var}[\operatorname{WFH}]} > \frac{0.18}{0.96} \approx 0.19,$$

This bound implies that WFH is associated with at least a $(\exp(0.19) - 1) \times 100 \approx$ 21 percent higher firm productivity level relative to non-WFH jobs. Since the estimates

¹⁹See also Aksoy et al. (2022).

²⁰Technically, λ is the slope conditional on WFH. The WFH share is small enough that omitting WFH in the AKM regression is likely not consequential.

on the causal impact of WFH on productivity are quite mixed and even the upper-end estimates are not that large, this likely reflects in large part an equilibrium relationship between WFH provision and unmeasured job characteristics.²¹ The findings imply that firms that offer WFH jobs tend to be more productive and recruit more skilled workers. But conditional on someone being in that type of firm, there is negative selection on skill into WFH roles. If unmeasured job productivity and rent-sharing is dominant, as it appears, it is understandable why even panel models have struggled to find sensible compensating differential relationships. Panel models control for individual skill, but not for firm-level productivity differences.

Adverse Selection The negative selection into WFH reported in Emanuel and Harrington (2024) raises the interesting possibility of a classic adverse selection problem in WFH. As Emanuel and Harrington (2024) note, if the preferences of WFH are negatively correlated with the unobserved productivity, there are conditions under which WFH will be under-provided, as the WFH wage has to be set so low that only people with the most extreme preferences will accept these contracts²² Whether adverse selection is a pervasive problem is an open question. On average, people with higher levels of education have higher levels of amenities, which points against the negative selection assumption. However, as we have seen, there are multiple levels of selection in terms of who sorts to high productivity firms versus who sorts into high amenity jobs conditional on firm productivity.

5.2.2 Injury and Fatality Risk

I now turn to the evidence on compensating differentials for injury and fatality risk. This has been a hugely influential area owing to the use of hedonic regression for job risk to calculate the VSL, and also because it speaks to the benefits of safety regulations.²³

²¹See Bloom et al. (2015) and Emanuel, Harrington, and Pallais (2023).

²²See also Nekoei (2022), who develops a model of adverse selection with amenities.

²³There has been considerable debate—particularly active in past decades—on whether workplace safety regulations are necessary if wages already reflect risk and workers have full information (see, e.g., Kniesner and Leeth 1995). This debate centers on efficiency and the possibility of market failures. For instance, imperfect information about workplace hazards can justify regulatory intervention, even though a robust tort system might otherwise incentivize employers to invest in safety. In the United States, however, workers' compensation typically substitutes for tort liability, and because it is only partially experience-rated, employers may still underinvest in safety.

Even if there is no market failure, distributional considerations can justify workplace safety regulation. Labor unions are among the primary stakeholders of the U.S. Occupational Safety and Health Administration (OSHA). Because union contracts frequently include health and safety provisions (Gray, Myers, and Myers 1998), mandatory standards codify protections that would otherwise require negotiation, freeing unions to bargain over other issues such as pay. Unions may push for safety provisions to ensure that wage gains in collective bargaining are not offset by worse working conditions, including safety. OSHA

This is also an area where the use of stated preference surveys are not very helpful since respondents may have a difficult time assessing low-probability events.

The Fatality Risk Premium and the Value of Statistical Life The value of a statistical life (VSL) is a central input in regulatory impact analysis, used to monetize the benefits from policy-induced mortality improvements (OMB 2003). It is also used widely for valuing health expenditures and other social programs (see, e.g., Finkelstein, Hendren, and Luttmer 2019). The basic idea stems from researchers trying to estimate an individual's MWTP for a reduction in mortality risk. In practice, the vast majority of VSL estimates come from hedonic wage regressions which relate variation in wages to variation in occupational fatality risks under the assumption that observed wage differentials reflect compensating differentials for risk exposure.

Let p_i represent the probability of a fatal accident in job *i*, measured as the expected number of deaths per 10,000 workers annually, and let w_i denote the corresponding wage. A standard specification regresses wages on fatality risk and other covariates:

$$w_i = \alpha + \beta p_i + X_i \gamma + \varepsilon_i,$$

where X_i is a vector of worker and job characteristics and β captures the wage-risk trade-off. Because one unit of p_i corresponds to a 1-in-10,000 increase in annual mortality risk, the implied VSL follows directly from the estimated β :

$$\text{VSL} = \frac{\Delta w}{\Delta p} = \beta \times 10,000.$$

This framework assumes that workers accurately perceive job-related fatality risks and that they sort into jobs such that observed wage differentials reflect an equilibrium premium for bearing additional risk. All of the points discussed earlier on the challenges of hedonic estimation apply here, particularly unobserved worker and job attributes. An example where it is simply impossible to disentangle attributes might be working a register in a night shift. A night schedule is harder to balance with family life, and it is more dangerous.²⁴

Despite the potential pitfalls, VSL estimates derived from these wage equations are commonly used by regulators. The U.S. Department of Transportation, for example,

also gives all workers a formal channel to exercise "voice"—for example, by filing formal grievances and complaints, but unions have historically been better able to use this mechanism (see, e.g., Mas 2008). Additionally, larger (and often more unionized) firms may prefer stringent regulations because they can more easily navigate compliance, thereby creating barriers to entry for smaller competitors.

²⁴Of course, job risk measures are not at that level of granularity.

uses VSL to monetize the benefits of safety regulations. The value they use—\$9.6 million in 2015—is derived from a meta-study that primarily uses studies relating wages to fatality risk. The US Environmental Protection Agency uses VSL to quantify the benefits of regulations affecting air quality, water safety, and hazardous waste management. As of 2024, it derived its VSL estimate from 17 hedonic wage analyses and five contingent valuation studies that were conducted between 1974 and 1991 and subsequently updated to reflect inflation and real income growth (Cropper, Joiner, and Krupnick 2024). There are many more such examples in the U.S. and globally, hence the policy relevance of this application of labor economics.

Measurement A concern specific to estimation of the workplace risk premium is the accurate measurement of fatality risk. Such a measure would ideally reflect a worker's understanding of risk, again assuming they are well-informed. If we assign the fatality risk at a broad industry level, if may not correspond to the risk in their actual job or establishment. However, using firm or establishment fatality rates has proven problematic because fatality is a rare event and such a measure would be noisy, resulting in estimates that are highly prone to attenuation bias. Researchers also face a methodological choice between using industry-level fatality rates or occupation-specific rates. In most datasets, industry-level measures generally have larger sample sizes and fewer small-cell estimation issues but may mask substantial heterogeneity in job tasks, safety standards, and worker characteristics. Occupation-specific measures can better capture the relevant risk exposures but often rely on smaller samples, which can increase sampling variance and measurement error.²⁵ But even an occupational measure will have error. The cashier position at a convenience store is very different than the at a high end boutique. As we have to rely on proxies for fatality risk, measurement error, both classical and non-classical, is a serious issue in estimation of compensating differentials for job risk.

Robustness and Publication Bias While many of the correlational studies that seek to estimate compensating differentials have counterintuitive signs or yield surprising magnitudes (e.g., the WFH studies summarized above), it is noteworthy that the VSL estimates used by regulators and derived from wage studies are fairly grouped together, point in the expected direction, and are reasonably close to VSL estimates from other settings like transportation, housing, and health studies. It is tempting to conclude that

²⁵The Census of Fatal Occupational Injuries (CFOI), compiled by the BLS, is often considered the most comprehensive data source for workplace deaths; it matches multiple data systems, such as workers' compensation files, death certificates, and administrative reports, to create a relatively complete record.

there is something unique about fatality risk that overcomes the estimation challenges we see with other work attributes. Unfortunately, this seems more a function of publication bias and methodological choices that are made on the basis of whether the resulting estimates are "reasonable." In a report commissioned by the EPA on VSL estimates, Black, Galdo, and Liu (2003) use three different datasets and a variety of specifications to assess robustness of VSL estimates using wage-risk data. Their primary conclusion is: "First, and foremost, the [VSL] estimates are quite unstable. Small changes in the specification of covariates or the risk measured used result in large variations in the estimated price of risk. Many of the estimates indicate that the price of risk is negative, which is contrary to the theoretical framework used." Others who have done this exercise have found the same (e.g., Hintermann, Alberini, and Markandya 2010).

Publication bias in VSL estimates appears severe, as found by Ashenfelter and Greenstone (2004) in the context of road safety. Viscusi (2015) concludes that publication bias in wage-risk studies is substantial, and can inflate estimates considerably. Viscusi argues that much of this bias stems from older studies that relied on voluntary reports of workplace fatalities rather than using the higher-quality Census of Fatal Occupational Injuries, which tends to produce higher estimates. Related to the concern of publication bias is the tendency for regulators to cherry-pick estimates from published results. An example of this can be seen in the U.S. Department of Transportation Revised Safety Guidance in 2016 that defines the VSL (USDOT 2016). Of the 15 published wage studies considered in their meta-study, two were dropped from the VSL guidance value due to "implausibly high" values.²⁶ There are clearly several layers of selection beginning from decisions on data and methodology to whether the study is published, to whether it is used by regulators. With many dozens of wage-risk estimates of varying quality, researchers and policymakers must ultimately determine which studies are highest-quality. Ideally, these judgments would not be made ex post, based on whether the estimates are deemed too large or too small. This is also a situation where pre-specifying data and models would be highly beneficial.

One recent observational study that takes seriously some of the crucial measurement and identification challenges is Lavetti and Schmutte (2018). They have two key advantages: first, employer-employee matched data, and second, high-quality administrative data that provides a complete census of worker fatalities. This allows them to measure occupational fatality risk with a level of precision previously unattainable. Lavetti and Schmutte (2018) use cross-establishment job changes and within-establishment varia-

²⁶Six studies were dropped in total for the following reasons: (1) "Implausibly high; industry-only risk measure," (2) "Occupation-only risk measure," (3) "Implausibly high; industry/occupation risk measure,"
(4) "Industry-only risk measure; no full-sample VSL estimate" (two studies), and (5) "VSL estimated only for occupational drivers."

tion in fatality rates by occupations to control for worker and establishment heterogeneity, and match effects. In doing so they can control for all stable establishment amenities that are correlated with both fatality rates and wages.

In their regression of log wages on fatality risk per 1,000 full-time full-year-equivalent workers with no additional controls they obtain an estimate of is -0.18. This negative relationship accords with the intuition that lower-skill workers sort into riskier jobs. After adding covariates, including educational attainment, occupation and industry controls, experience, and race the coefficient increases to 0.279 and has the "expected" sign in the sense that higher wages are required to compensate for added risk. Adding worker effects reduces the estimate to 0.037. As discussed by Lavetti (2023), the attenuation in the risk variable when adding individual fixed effects is observed in other wage-risk papers, for example Kniesner et al. (2012). It is a puzzle because the addition of education controls confirms that job risk is negatively correlated with skill (the estimate increases from -0.18 to 0.279), yet adding individual fixed effects which absorbs both observed and unobserved skill reduces rather than further increasing the risk premium.²⁷ One explanation for this reversal is that attenuation bias due to classical measurement error can be more severe in fixed-effect models.²⁸ Given the high-quality fatality measure, a more likely reason for this reversal, as suggested by the authors, is that there is a job ladder, and jobs at the bottom of the ladder are unsafe, unpleasant, and low-pay. Mobility up the job ladder is associated with decreases in fatality risk along with increases in compensation and other improvements that are hard to measure. If the model is identified off of these moves it will downwardly bias the risk premium. In Lavetti and Schmutte (2018)'s preferred specification, they control for worker fixed effects, match effects, and employer effects and obtain an estimate of 0.17. They conclude that unobserved establishment effects are important for the estimated risk premium.

The main take-away from both the WFH and the workplace safety literature, and indeed from most other domains, is that there is variation in firm productivity, and more productive firms offer both higher wages, more generous amenities, and better working conditions. Workers with greater earnings potential are more likely to be employed in these firms at higher rungs in the job ladder.

²⁷The OVB formula is Bias = $B_{\text{estimated}} - B_{\text{true}} = B_{\text{skill}} \cdot \frac{\text{Cov(fatality risk, unmeasured skill)}}{\text{Var(fatality risk)}}$. If $B_{\text{skill}} > 0$ and Cov(fatality risk, unmeasured skill) < 0, then the OVB bias is negative.

²⁸To see this, let amenity A_{jt} follow an AR(1) process: $A_{jt} = Z_{jt} + e_{it}$, where $e_{it} = \rho e_{i,t-1} + s_{it}$, $|\rho| < 1$. In a cross-section, regressing $w_{ijt} = B_t Z_{jt} + \mu_{ijt}$ on A_{jt} yields $B_m^{\text{CS}} = B_t \frac{\operatorname{Var}(Z_{jt})}{\operatorname{Var}(Z_{jt}) + \operatorname{Var}(e_{it})}$. With first-differences, $\Delta w_i = B\Delta Z_j + \Delta \mu_i$ and $\Delta A_j = \Delta Z_j + \Delta e_i$ imply $B_m^{\text{FD}} = B \frac{\operatorname{Var}(\Delta Z_j)}{\operatorname{Var}(\Delta Z_j) + \operatorname{Var}(\Delta e_i)}$. Under stationarity, $\operatorname{Var}(e_{it}) = \sigma_s^2/(1-\rho^2)$ and $\operatorname{Var}(\Delta e_i) = 2\sigma_s^2/(1+\rho)$, which can be smaller or larger than $\operatorname{Var}(e_{it})$ depending on ρ .

Experimental and Quasi-Experimental Approaches In principle, experiments or quasi-experiments that induce plausibly exogenous changes in amenity provision or costs can overcome many, though not all, of the obstacles discussed above. The ideal experiment would be to observe a set of firms making an exogenous change to an amenity, and then let enough time pass that this change can be reflected in firm outcomes, including wages. Unfortunately, such designs remain rare. Here I discuss several of these studies, again in the context of workplace health and safety.

A good example of an exogenous cost shifter are randomized safety inspections. If they reveal violations, the firm has to either pay a fine or correct the deficiency. This contingent fine lowers the marginal cost to the firm for making the safety improvement. Levine, Toeffel, and Johnson. (2012), Johnson, Levine, and Toffel (2023) and Lee and Taylor (2019) use data on randomized safety inspections to learn about their impacts on injuries, employment and payroll. Levine, Toeffel, and Johnson. (2012) estimate a 9% decline in injury counts and a 30% decline in worker's compensation claims following randomized inspections in the previous four years by California's OSHA equivalent. The worker's compensation administrative records they use also contains employer payroll, and here they do not find a significant impact, with a the 95% confidence interval for the treatment effect of (-2.0%, 3.0%). By contrast, Lee and Taylor (2019), who use data from the Census of Manufacturers, estimate a statistically significant 2-3% reduction in wages, computed as the total annual payroll for production workers in a plant divided by total hours worked by production workers, following an inspection. Whether this wage reduction is interpretable as a compensating differential requires ruling out rent-sharing by assumption. If correcting deficiencies were costly or lowered productivity, with rent-sharing that would translate to a lower wage. As we have considerable evidence of rent-sharing, that is a strong possibility.

Another paper that uses experimental variation is Boudreau (2024), who implemented a field experiment with multinational apparel companies to randomly mandate the establishment of Occupational Safety and Health committees at about half of the 84 Bangladeshi suppliers for the companies. These committees facilitate communication and coordination between workers and managers to ensure safe working processes. These types of committees are common in collective bargaining agreements globally, including in the U.S. (Gray, Myers, and Myers 1998). The short-run findings are that these committees improved safety compliance, but did not have measurable impacts on injuries, wages, or employment. Interestingly, they did have significant negative impacts on worker satisfaction. This last finding is an additional reminder that amenities often do not operate in the workplace as expected. In this case, it is possible that the low adoption of these committees absent the mandate was because the marginal valuation of safety compliance among workers did not compensate for the relative inconvenience.

An example of a quasi-experiment that dramatically changed working conditions is Wissmann (2022), who studies an indoor smoking ban in Germany. He focuses on "mini-jobs" in the hospitality sector, which include marginally-employed persons, typically in temporary or part-time employment arrangements, with capped compensation. By focusing on these mini-jobs he has as close as possible to a spot market, with lower wage rigidities and more limited rent-sharing. He finds that after the smoking ban daily earnings of workers in bars and restaurants fell by 2.5%. This is one of the rare papers with quasi-experimental identification that shows a result that is directionally consistent with the theory of compensating differentials.

6 Market-level Incidence and Mandates

The positive analysis of amenities is informative about the structure of markets and firm and worker preferences. But much of the debate around amenities is policy-oriented and market-level and deals with the welfare and distributional effects of policies that mandate or restrict amenities. In this section, I discuss the economic incidence of market-level changes in costs to employer-provided benefits, market-level amenities and mandates. Classic references are Roback (1982) and Summers (1989). A distinction between this and the compensating differentials exercise we have already covered is that at the broader market or industry level, changes in the cost of benefit provision can have a number of equilibrium effects that ultimately have ambiguous impacts on the wage.

Labor Market Impacts of Rising Health Costs An important question that I will use to illustrate the main elements of this area of inquiry is: What are the labor market effects of rising health costs? The seminal paper is Gruber and Krueger (1991), who look at the worker's compensation program, an insurance program administered by U.S. states that provides wage replacement and health benefits for workers injured on the job. Employers pay a premium that is partially experience-rated, and the benefit is mandated, so any increase in health care costs will translate into increased employer costs through a higher premium. In that case, Gruber and Krueger show that wages will fall according to

$$\frac{dw}{dp} = -\frac{\eta^D - \alpha \eta^S}{\eta^D - \eta^S},$$

where p is the employer-paid insurance premium, η^D is the elasticity of labor demand, η^S is the elasticity of labor supply, and $\alpha \times p$ is the employee's monetary value of the

benefit. As in Summers (1989), if workers fully value the benefit ($\alpha = 1$), the change in the wage is just a compensating differential, with all increases in cost shifted to lower wages. If $\alpha = 0$, workers do not value the benefit and the wage response is governed by the standard tax incidence formula.²⁹

A related question is how wages and employment respond to benefit costs when benefits are not mandated and employers can decide on generosity levels. This generosity margin appears potentially empirically important as the aggregate trend has been towards a lower employer share of premiums, higher cost-sharing, and a lower share of firms offering health benefits (Claxton et al. 2024).

To see how changes in generosity margin interact with the response to wages, consider a representative firm that provides total compensation T to a worker through a combination of wage w and benefits b. Let the worker's total compensation be T = w + g(b), where g(b) is the monetary value to the worker of the benefits and g' > 0and g'' < 0. The firm faces a per-worker cost of w + cb, with c > 0 denoting the perunit price of benefits. Although T necessarily declines when c increases, the effect on the wage w = T - g(b) is ambiguous, because the benefits component b also changes endogenously. The wage responds to benefit cost according to:

$$\frac{dw}{dc} = \frac{dT}{dc} - \frac{d}{dc} \left[g \left(b^*(c) \right) \right] = \frac{dT}{dc} - c \frac{db^*(c)}{dc}$$

The first term is negative. The second term, $-c \frac{db^*(c)}{dc}$, is positive because $\frac{db^*(c)}{dc} = \frac{1}{g''(b^*(c))} < 0.^{30}$ The degree to which firms cut back generosity is a function of diminishing marginal utility of benefits. Benefits are cut more if small changes in benefits have little impact on utility. If *b* adjusts strongly to the increase in costs, the wage gain from substituting away from benefits can outweigh the drop in *T*. Otherwise, the wage falls. Ultimately the labor market impacts of changing benefit costs is an empirical question. As I discuss, a reasonable characterization in the literature is that changes in health costs have employment effects, limited effects on benefits (with the caveat that these are difficult to measure at a granular level), and mixed evidence on pass-through to wages.

Baicker and Chandra (2006) use differential growth in malpractice costs across states as a (somewhat weak) IV for premiums. They find that a 10 percent increase in premiums is associated with a 1.2 percentage point reduction in the aggregate prob-

²⁹An interesting question is whether in the case of worker's compensation employers also adjust on the margin of investing in safer workplaces. I am not aware of a study that tests this channel using modern empirical methods, but it is a plausible channel for the significant declines in workplace injuries over the last decades.

³⁰The benefit $b^*(c)$ is chosen to satisfy $g'(b^*(c)) = c$.

ability of being employed. Their estimates on the provision of health insurance are imprecise, but arguably these measures are too coarse since the relevant margin of adjustment for employers is not the provision but the generosity of health insurance in terms of contributions. Their estimates on wages for workers with employer health insurance suggest a dollar-for-dollar reduction in wages on rising premiums.

A number of papers look at labor market impacts from hospital mergers. There has been significant consolidation of hospitals in the U.S., and there is consistent evidence that this has led to rising prices charged to insurers. For example, Brot-Goldberg et al. (2024) find that hospital mergers in already concentrated markets raised prices by 6%. Cross-market mergers result in increases in the range of 6-12% (Dafny, Ho, and Lee 2019; Arnold et al. 2025). These price increases translate to higher health care costs to employers in these markets via higher premiums. Using a sample of mergers, Brot-Goldberg et al. (2024) document that a 1% increase in health care prices induced by mergers lead to a 0.27% drop in income at the county-level. This reduction in income comes from reduced employment rather than lower wages.³¹ In the above model that result would imply that benefit generosity is the margin of adjustment, but they do not find evidence that employers shift to high-deductible health plans.

Brot-Goldberg et al. (2024) provide evidence in support of the head-tax interpretation of employer-sponsored health insurance, that is, the idea that there is a fixed cost for low-wage workers due to the provision of benefits that are optimized for higher-wage workers and, hence, valued by low-wage workers at less than cost. This fixed-cost is hypothesized to lead to low joint surplus and more sensitivity to shocks. Brot-Goldberg et al. (2024) find that hospital mergers do not cause unemployment for workers at the top of the income distribution. Workers at the very bottom of the distribution, those earning less than \$20,000, also did not experience elevated unemployment, presumably because they have very low baseline rates of health insurance coverage and so experienced no shock. Unemployment effects were concentrated in the group of low-income workers who likely had health insurance. This evidence is consistent with these lower income workers in jobs with lower initial joint surplus, which was eliminated after the cost shock. What is less clear is whether this lower surplus was due to the health insurance provision and the associated head-tax wedge, or just a normal gradient of the joint surplus as a function of the productivity of the match.

Inelastic Factors in the Worker Production Function The discussion on benefit generosity in response to cost shocks makes clear that wages alone cannot be used to infer preferences over amenities when defined at the market-level. We also have to keep

³¹See also Arnold and Whaley (2020).

track of quantities. In the case where there are inelastic factors required for worker entry, welfare calculations also require accounting for price changes across all factors. The classic example is when firms have locational amenities, as in the Roback (1982) model, and workers have to rent housing. But this inelastic factor can take other forms, such as childcare or training and certification. If a worker requires childcare to work, and childcare is inelastically supplied, by the Roback logic the incidence of a marketchange in amenity rents will be on the childcare providers and MWTP= $q_c \frac{dp_c}{da} - \frac{dw}{da}$, where q_c is the quantity of childcare purchased by workers and p_c is the unit price childcare. The intuition for this finding is that if there is a small increase in firm amenities, perhaps through a policy measure, working becomes more attractive, labor supply shifts up, and price of child care providers are bid up. In this case, absent data on the quantity and price of this inelastic factor, we are not able to use compensating differentials to measure preferences. More generally, the full impact of a change in amenities in a market requires measuring a weighted sum of the change in price or shadow price for all factors that are required for worker entry.

7 The Analysis of Non-Wage Amenities Under Imperfect Competition

The competitive model is useful as a benchmark but, as has been emphasized, any understanding of the provision, valuation and implication of non-wage amenities has to be considered in more realistic environments that are imperfectly competitive. Imperfect competition affects the provision of amenities such that it is no longer viable to assume that amenities are provided to always achieve a market reservation utility. As predicted by the representative agent model outlined in this Chapter, as long as amenities are normal goods, any increase in the value of a job (the \overline{V} in problem (1)) is achieved by increasing wages and amenities. Therefore, variation in job value will generally result in positive co-movement in wages and amenity provisions. Variation in job value implies job ladders, and higher rungs of the ladder are better in more than one dimension. There are good and bad jobs.

Equilibrium variation in job value arises in a general class of models with imperfect competition. For example, if a firm or industry has rents, e.g. due to imperfect product market competition, it may share those rents with workers. This may be because of the sociology of organizations ("Wages over a considerable range reflect managerial discretion. When management can easily afford to pay high wages, they tend to do so" (Slichter 1950)), or because a profit-maximizing firm finds it optimal to raise

wages above competitive levels to reduce turnover (Salop 1979), or as a substitute to monitoring (Shapiro and Stiglitz 1984), or for morale considerations (Akerlof 1984).

An interesting feature of models with endogenous amenity determination is that the rents in question can arise internally in the model. For example, in the Rosen model I outlined, firm rents are increasing in type, that is, the lower the cost of amenity provision the higher is the level of rents. To the extent that there is rent-sharing, this will shift the hedonic wage function since in Figure 3 type-3 firms will offer a $(a_3^*, w_3^* + r \times \pi_3)$ bundle, where *r* is the rent-sharing parameter. Since profits π_k are growing in type, the slope of the hedonic wage function will be attenuated.

Rents also arise naturally in models of amenity determination with search frictions. Key early references here are Gronberg and Reed (1994), Hwang, Mortensen, and Reed (1998), and Lang and Majumdar (2004). In essence, these models show that with search frictions there is equilibrium dispersion in job value and that an endogenous relationship emerges between firms that provide more amenities and firm surplus. Wages are downstream of surplus, thereby breaking the direct link between amenities and wages emphasized both in the Rosen model and in hedonic models. I will review the mechanics of a search model with a slight variant of Hwang, Mortensen, and Reed (1998).

7.1 Amenity Determination with Search Frictions

The environment of Hwang, Mortensen, and Reed (1998) is as follows. Jobs consist of a wage rate, w, and a non-wage amenity, a. Workers are homogeneous in productivity and preferences. Firms are heterogeneous in their cost efficiency of providing a. Denote these costs c_j . Workers search for job opportunities, moving from unemployment to employment and between firms. The probability of a job offer is given by λ_U and λ_E for workers in unemployment and employment, respectively. The separation rate into unemployment is given by δ . The distribution of job values is given by $\overline{V}_j \sim F(\cdot)$, where \overline{V}_j is the job utility offered by a specific firm. The distribution $F(\cdot)$ is endogenously determined in the model. The utility derived from a job with attributes (w, a) is given by:

$$V(w,a) = w + h(a),$$

where h(a) is a quasi-concave function representing the utility derived from the amenity. Workers' MWTP for *a* is given by h'(a).

Firms choose (w, a) to maximize profits:

$$\pi_j = [p - w - c_j a] m(\overline{V}), \tag{6}$$

where $c_j(a)$ is the per-worker cost function for type-*j* firms, and *p* is revenue per worker. The equilibrium firm size function is $m(\overline{V})$ and is shown to be increasing in \overline{V} . Hwang, Mortensen, and Reed (1998) show that in equilibrium firms have positive profits.

In the model the amenity and wage mix are chosen efficiently in two steps. First, amenities are chosen efficiently to equate MWTP to MC following $h'(a_j) = c_j$. Then, wages are set to match total job value. Classic "Rosen" compensating differentials emerge in the form of an inverse relationship between wages and amenity values for a fixed job value. Variation in job value obscures these equalizing differences.

There are two reasons job values vary in this model. The first is due solely to search frictions. As is standard in search models with posting, there is equilibrium wage dispersion. To isolate this mechanism we shut down heterogeneity in the cost of amenity provision so that all firms are identical. Then we invoke the usual Burdett and Mortensen (1998) argument that without job value dispersion a firm can incrementally increase its wage offer and obtain a discontinuous increase in employment, and thus revenue, with only a negligible change in labor costs. The resulting increase in profit is in contradiction with equilibrium. Therefore, in equilibrium there must be dispersion in job value for otherwise identical firms. Because amenities are pinned down by the first-order condition, job value and wages move in tandem.

The second reason job values vary in Hwang, Mortensen, and Reed (1998) is due to heterogeneity in cost of amenity provision. Search plays only an indirect role in this channel, through the size function $m(\overline{V})$, which depends on \overline{V} because of frictions. Framed this way, it is the static monopsony problem. To proceed, and to isolate the role of cost variation, I will consider the static case below.

7.1.1 Static Exposition as a Monopsony Problem

In the static problem the firm maximizes profits given by equation (6) where, as before, utility is V = w + h(a). The function $m(\cdot)$ is increasing and concave (i.e., m'(V) > 0 and m''(V) < 0), and $h(\cdot)$ is also increasing and concave (i.e., h'(a) > 0 and h''(a) < 0). Thus, the firm's problem is $\max_{w,a} (p - w - ca)m(w + h(a))$. As before, the optimal amenity is pinned down by the first-order condition, where $a^*(c)$ satisfies $h'(a^*(c)) = c$. Since h''(a) < 0, it follows that optimal amenity provision is decreasing in cost: $\frac{da^*(c)}{dc} < 0$. The optimal wage choice solves $p - w - ca = \frac{m(V)}{m'(V)}$, and because V = w + h(a), this becomes $V + \frac{m(V)}{m'(V)} = p - ca + h(a)$. Defining $R(c) \equiv p - ca(c) + h(a(c))$, we see that the chosen V must satisfy $V + \frac{m(V)}{m'(V)} = R(c)$. Defining, $F(c,V) \equiv V + \frac{m(V)}{m'(V)}$ by the Implicit Function Theorem, we obtain $\frac{dV}{dc} = -\frac{\frac{\partial F}{\partial c}}{\frac{\partial F}{\partial V}} = -\frac{a(c)}{1 + [m'(V)^2 - m(V)m''(V)]/[m'(V)]^2} < 0$

0. Hence, as the amenity cost c increases, the firm's optimal total compensation V decreases.

This result shows that when the labor supply function to the firm is upward sloping with respect to job value, variation in amenity cost leads to variation in target utility, lower cost firms offer higher job values. This is the static-version of the result in proposition 5 in Hwang, Mortensen, and Reed (1998). This result shows that lower amenity costs are a source of rents that lead to differences in job value. Another implication of this model is that amenities reduce markdowns and raise employment closer to competitive levels. The reason is that an efficiently provided amenity is pinned down by the first-order condition, and the firm takes this as a given prior to the firm setting the wage. The amenity shifts the labor supply curve outward, thus increasing employment and lowering the markdown.

7.2 Horizontal Amenities and Thinness in the Distribution of Preferences

Non-wage amenities are the main ingredient of imperfect competition arising due to "thinness" of preferences in the labor markets. This concept relates to the notion of horizontal amenities, over which individuals have idiosyncratic preferences. Commuting is the prototypical example of this phenomenon, as discussed by Manning (2003). Places that are nice to live in and places that are productive for firms do not necessarily overlap. Workers make choices on where to live and work, balancing housing costs, wages and commute times. In a frictionless competitive market, all workers in a firm will have identical preferences over residential location, by the Rosen logic. The set of workers with this preference is thick enough that the firm takes their preferences as given and can hire any number of workers at the prevailing market wage. But if the distribution of workers willing to pay the commuting cost is thin, the firm may "move through" this distribution of preferences, and the conditions are set for monopsonytype behavior. The resulting upward-sloping labor supply curve to the firm represents the compensating differential for the marginal worker over the distribution of preferences. As has been shown by Card et al. (2018), we then obtain the standard result of a markdown relative to marginal product, which is often expressed as a weighted average of marginal product and an outside option where the weights are a function of the labor supply elasticity to the firm. In this example, underlying the question of whether the firm is a price-taker is whether the firm views the average preferences of its workforce as a quantity that can be influenced by its hiring actions.

This framework is the product differentiation version of monopsony that is out-

lined in Card et al. (2018) and can arise whenever there is a horizontal amenity with a sufficiently thin distribution of preferences facing the firm. A horizontal amenity is anything that gives a worker a match-specific utility with a firm.³² Despite horizontal amenities increasingly being used as a modeling technique to generate firm markdowns, we could use more direct evidence on what differentiates firms. The commuting cost example is more like a classic monopsonist because it requires a firm that can hire a large share of its workforce within its close geographic proximity. But one can think of other non-pecuniary amenities that are not easily replicated and over which workers have heterogeneous preferences, for example the presence of friends at the workplace or a company's culture, which can be deeply rooted in founding principles.³³

If the source of market power is the firm as a differentiated product, what are the welfare implications? The social cost is from unrealized profitable production opportunities as these firms are inefficiently small. Worker welfare is less clear. Horizontal amenities that on the margin give weakly positive utility are, yet again, Ricardian rents that benefit workers. Worker welfare is higher than if these amenities were shut off and their resulting market power effects were eliminated. This is because inframarginal workers are strictly better off than in their next alternative, even with a markdown. Friends at work gives workers surplus even if their wage is lower due to the compensating differential. The markdown shades that surplus to some degree but does not eliminate it. If the source of market power is differential commute times, the welfare consequences are less clear and would require a model with residential and firm location decision, in the Roback spirit, to evaluate. In the commuting case we are closer to the classic monopsony case. Labor markets in this case are not exactly company towns in the Joan Robinson sense, but they are a smoothed version of that scenario because a thin distribution of preferences requires a degree of geographic isolation. Even then, markdowns require special circumstances. A large employer in a low density geographic area is capable of moving through the distribution of commuting preferences in a local area. But if the marginal hire is in a distant city that has elastic supply, there will be no markdowns, even if the firm is isolated locally.

Idiosyncratic preferences over firm amenities have become a widely-employed modeling tool to generate mark-downs in wages. However, a tension in these models, and also monopsony models microfounded on search frictions, is that monopsony can be unwound if there is also heterogeneity in skill. If skill is heterogeneous, a firm has to balance its ability to markdown wages due to horizontal amenities with changing skill

³²Any vertical amenity in which worker's have heterogeneous preferences can be modeled as a combination of a vertical and horizontal amenity. This would be done by treating the mean-zero heterogeneous term as horizontal and amenity level as vertical.

³³Park (2016) estimates the median WTP to work with a friend at 4.5% of wages.

composition if it does so, because more productive workers have better outside options. A firm facing both idiosyncratic preference and skill heterogeneity will be constrained from marking down wages because the marginal revenue = marginal cost optimization equality is premised on having a level of worker skill. If average skill of the workforce is a function of the wage, wage markdowns will decrease average skill, thus breaking the optimization condition.

This role of skill heterogeneity can be seen in a simple model. Suppose workers are characterized by two independently distributed attributes: skill s_i and idiosyncratic preferences over the firm amenities θ_i . A worker accepts a job (offered at wage w) if and only if $w + \theta_i \ge s_i$, where s_i is the productivity of the worker in the competitive market and represents their outside option. Skill enters the model through an elevated outside option.³⁴

The firm's revenue from hiring a worker of skill *s* is $p \times s$, where *p* is output price. The firm's profit when the wage is set at *w* is

$$\pi(w) = p\left[S(w) \times L(w)\right] - w\left[L(w)\right], \tag{7}$$

where $L(w) = \Pr(w + \theta_i \ge s_i)$ is the share of workers who accept the firm's offer at w and $S(w) = \mathbb{E}[s_i|w]$ is the expected skill among those who accept, conditional on the wage. By the choice problem, both of these functions are increasing in w. Average skill will be more sensitive to the wage when the variance of θ_i is lower. When idiosyncratic preferences have higher variance, there will be more situations where a high-skill worker is willing to work at a firm for a low wage. Define elasticities $\varepsilon_L(w) = \frac{d \ln L(w)}{d \ln w}$ and $\varepsilon_S(w) = \frac{d \ln S(w)}{d \ln w}$. Because revenue is p[L(w)S(w)], the markdown is:

$$pS(w^*) - w^* = pS(w^*) \left[1 - \frac{w^*}{pS(w^*)} \right] = pS(w^*) \frac{1 - \varepsilon_S(w^*)}{1 + \varepsilon_L(w^*)}.$$
 (8)

If $\varepsilon_S(w^*) = 0$ (raising *w* does not upgrade skill composition), we revert to the classic monopsony wedge $= \frac{pS(w^*)}{1 + \varepsilon_L(w^*)}$. If $\varepsilon_S(w^*)$ is large, the markdown decreases.³⁵ In the limit $\varepsilon_S \to 1$, it can vanish. An implication of this model is that markdowns are more sustainable when markets are relatively heterogeneous in preferences over firm-specific amenities but homogeneous in skill. The elasticity of average productivity with respect

³⁴See also Kline (2024) for a related screening model that combines skill and monopsony power.

³⁵Derivation of the markdown: Differentiating equation (7) and setting $\frac{d\pi}{dw} = 0$ yields p[L'(w)S(w) + L(w)S'(w)] = L(w) + wL'(w). Dividing both sides by pL(w) and using elasticity definitions gives $\frac{S(w)}{w} [\varepsilon_L(w) + \varepsilon_S(w)] = \frac{1+\varepsilon_L(w)}{p}$. Rearranging gives $\frac{pS(w)}{w} = \frac{1+\varepsilon_L(w)}{\varepsilon_L(w)+\varepsilon_S(w)}$, so $\frac{pS(w)-w}{pS(w)} = \frac{1-\varepsilon_S(w)}{1+\varepsilon_L(w)}$. Multiplying through by pS(w) then establishes formula (8).

to the wage is a quantifiable object that has been estimated in a number of contexts, such as in the public service in Mexico (Dal Bó, Finan, and Rossi 2013) and in U.S. companies (Emanuel and Harrington 2022). This quantity will be market- and context-dependent, but in the Emanuel and Harrington case, they estimate an elasticity close to 1, implying that skill sensitivity offsets any potential market power the firm they study might have due to horizontal amenities.

8 Revealed Preferences Approaches

In the standard hedonic framework, workers are assumed to choose job attributes as if selecting from a fully articulated menu. In practice, however, several complicating factors undermine this idealized setting. First, jobs often function more like experience goods than search goods, with their true value becoming apparent only after some period of employment. Second, market frictions and informational asymmetries limit the set of firms or positions that workers consider, effectively constraining their choices at any given time. Consequently, workers are frequently misallocated and must adjust their positions dynamically as they acquire new information and seek better matches.

When there are frictions, and the wage is not fully allocative, market clearing shifts to the quantity side. People move as they find better opportunities. Higher value firms will have fewer vacancies, less turnover, and better ability to scale. All of these features provide us with useful information. While imperfect and incomplete matching might complicate hedonic estimation, it presents opportunities because job choices and mobility can be used to estimate desired quantities.

In this section I will review the range of approaches that have built on the choice margin. The ways of doing so are varied. We can look at choice between different work arrangements at varying wages, choice between jobs and outside options at varying amenity (e.g., risk) levels, choice between multiple heterogeneous jobs, job separation decisions, and worker flows more generally.

Overall, a recognition that wages and other job characteristics are not perfectly equalizing has led to a successful methodological and empirical literature that has leveraged these market imperfections to estimate preferences over amenities in ways that the hedonic approach has not done consistently.

I organize this section as follows. (1) I begin with the simplest possible example, in which an individual chooses between two positions that differ along a single dimension. Drawing on Mas and Pallais (2017), I consider how to identify the demand function when we can randomly vary the relative wage. (2) Next, I expand the discussion to choice over job alternatives more generally, highlighting the identification assumptions

required when random variation is not available. I also discuss quasi-experimental approaches. (3) I then move to the search literature, outlining the approach of Gronberg and Reed (1994), who use hazards or job durations to infer MWTP. (4) Finally, I discuss how ranking algorithms, combined with data on worker flows, can be used to estimate the relative value of firms, following Sorkin (2018), and applications of this approach. Each of these approaches has different data and identification requirements: some require experiments, others can be implemented with survey data, and still others require employer-employee matched data. As we will also see, drawing on the discussions of Greenberg et al. (2021) and Anelli and Koenig (2021), having access to institutional features can be highly beneficial for identification in revealed preference models.

8.1 Discrete Choice Experiments

Consider the discrete version of the Rosen model in which a worker *i* chooses between two jobs, indexed by j = 0, 1. Job 0 offers a baseline amenity a_0 (e.g., a standard 9–5 in-office schedule) and wage w_0 , while job 1 offers an alternative amenity a_1 (e.g., work-from-home) and wage w_1 . Let

$$\Delta w = w_0 - w_1,$$

so Δw is the wage premium of the baseline job over the alternative. Let WTP_i be worker *i*'s WTP for a_1 . Then the worker chooses job 1 if and only if

$$WTP_i > \Delta w$$

If WTP_{*i*} follows a distribution $F(\cdot)$ with mean μ and variance σ^2 , the probability of choosing the amenity job is

$$\Pr(\text{Choose } a_1 | \Delta w) = \Pr(\text{WTP}_i > \Delta w) = 1 - F(\Delta w).$$

Varying Δw and observing how many workers choose a_1 can thus reveal $F(\cdot)$. Note that the fraction of workers who choose the amenity is $Q = 1 - F(\Delta w)$, so the population inverse demand for the amenity is $\Delta w = F^{-1}(1-Q)$. Once estimated, we have information about preferences of inframarginal workers, allowing us to do welfare calculations.

Mas and Pallais (2017) embed this choice problem in an actual application process for a call center. Their goal is to quantify worker valuation over a variety of alternative work arrangements, including the option to work from home, flexible scheduling, and irregular employer-driven scheduling. In this discrete choice experiment they offer varying alternative work arrangements at randomized relative wages. Job seekers could choose between a standard 9-5 office position and alternatives. These alternative arrangements were assigned random wages in relation to the 9-5 position.

As in the discrete Rosen model, in their setup, they compare two positions at different relative wages, where the difference in the wage between the standard 9-5 job and the job with an alternative work arrangement is Δw . As above, a job seeker accepts the job with an amenity if their WTP for the amenity exceeds the price which in this case is given by the difference wage. The probability of that occurring is $Pr(WTP_i > \Delta w)$, where WTP_i is allowed to be heterogeneously distributed with mean μ and standard deviation σ . Assuming normality, one can estimate a probit $Pr(WTP_i)$ Δw) = F($b\Delta w + c; \mu, \sigma$), where $F(\cdot)$ is the normal cdf, b is the slope coefficient on the relative wage, and c is the constant.³⁶ Typically in a discrete choice model, it is not possible to identify the standard deviation σ of the latent variable due to scale invariance. However, because Δw is in the same unit as the latent variable, and there are no other factors that influence choice, it is possible in this case.³⁷ Specifically, $\mu = -(c/b)$ and $\sigma = -(1/b)$.³⁸ Once estimated, this gives the entire distribution of WTP as well as the population inverse demand function for the amenity. As a sampling of the findings in Mas and Pallais (2017), the mean WTP as a percentage of wages, in relation to standard 9am-5pm Monday-Friday in-office jobs are: 8% for WFH, 0% for an employee-determined flexible schedule, -19% for weekend work, and -20% for an employer-determined schedule.

Once we have an estimate of the entire WTP distribution, or equivalently the inverse demand function, there is much we can learn, especially when combined with external information. Mas and Pallais use the estimated WTP distribution combined with the Rosen (1986) model to estimate MWTP and marginal costs of amenity provisions at observed levels in the economy. This is done by looking at actual shares of the work arrangement (e.g., WFH) in survey data, and then reading the implied WTP corre-

³⁶Mas and Pallais (2017) estimate a logistic, but I summarize it here using a probit for expositional purposes. They also measure inattention in the job choices and estimate a mixture model that controls for inattention rates to obtain unbiased estimates of b and c. They show that inattention does not bias the estimate of the mean WTP but does bias the variance.

³⁷Specifically, in a standard binary choice model, one posits y = 1 if $bx + \varepsilon > 0$, $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. Because only the sign of $bx + \varepsilon$ is observed, simultaneously scaling *b* and ε by any $\lambda > 0$ leaves the sign unchanged, thus precluding separate identification of σ . By contrast, in a threshold framework with y = 1 if WTP > Δw , where WTP ~ $\mathcal{N}(\mu, \sigma^2)$ and Δw is observed, one cannot freely rescale WTP or its variance without altering how P(y = 1) changes as Δw varies. This breaks the usual scale invariance and makes it possible to identify σ .

³⁸These come from: Pr(WTP_i > Δw) = 1 - $\Phi((\Delta w - \mu)/\sigma) = \Phi((\mu - \Delta w)/\sigma)$. Noting that $(\mu - \Delta w)/\sigma = (\mu/\sigma) - (1/\sigma)\Delta w$, we match this expression to the empirical probit form $\Phi(b\Delta w + c)$ to identify $b = -1/\sigma$ and $c = \mu/\sigma$. Solving these relationships gives $\sigma = -1/b$ and $\mu = -c/b$.

sponding to that share in the cdf of the estimated WTP distribution. For example, they estimate a WTP for the option to work from home of 8 percent of wages for the average worker.³⁹ But because the share of WFH was only 10 percent (at the time of the paper), the marginal worker valued it at 21 percent of wages. In a Rosen model the MWTP is equal to the marginal cost of provision, suggesting it would cost the marginal employer 21 percent of wages to implement this arrangement. Because recent estimates have not found a change in the WTP to work from home, variation over time should reflect a change in the relative cost of provision. During the Covid-19 pandemic the relative cost of working in the office rose dramatically, which led to large increases in the WFH rate. Subsequently, the cost declined and we saw relative reductions in WFH rates.

8.1.1 Survey-Based Discrete Choice Experiments

One of the most important recent developments for our substantive understanding of labor markets has been the growing use and acceptance of survey-based Discrete Choice Experiments (DCEs). In these surveys, participants are presented with hypothetical scenarios in order to gauge their preferences and estimate the economic value they place on certain attributes or outcomes. These are in the family of conjoint analysis which has long been used in marketing for pricing analysis. Studies that use these surveys, of which there are now many, have given us a direct way to measure and quantify people's preferences over features of the workplace.

DCEs typically involve showing respondents side-by-side comparisons of hypothetical jobs, consisting of their attributes. The attribute vector $X = \{w, a_1, ..., a_n\}$ randomly varies. These comparisons can be done as a table that lists the attributes, or in the form of a narrative vignette. These variations can be shown once or many times per respondent. The choice problem is conventionally framed as a random utility model and estimated with a conditional or multinomial logit. If there are multiple responses per respondent, a mixed logit that allows for preference heterogeneity is feasible, as employed by Maestas et al. (2023), for example. Hierarchical Bayesian methods can also be used to obtain individual level estimates. For high-dimensional attributes, using an adaptive method like the Bayesian Adaptive Choice approach in Drake et al. (2024) is helpful for statistical power.

While DCEs have long been used in health and transportation economics, for example to value public goods, they have been rarely used in labor economics until relatively recently. Early examples of DCEs in workplace settings are in the healthcare worker

³⁹A similar magnitude has also been found in other studies (e.g., Aksoy et al. 2022; De Fraja et al. 2024; Nagler, Rincke, and Winkler 2024)

context (Mandeville, Lagarde, and Hanson 2014) and to estimate fringe benefits (Eriksson and Kristensen 2014). While it is much easier to estimate preferences with a DCE than other alternatives, such as market-based data, the concern is that responses to hypothetical scenarios may not correspond to what people do in the real world. This, I suspect, is one reason labor economists have tended to prefer to look at what people do rather than what they say.

The widespread recent use of DCEs is the result of a recognition that the answers to these surveys align well with the ground truth when validated. Mas and Pallais (2017) show that estimates from DCEs embedded in actual applications result in WTP measures that are very close to survey-based hypothetical choices. For example, they estimate that real job applicants are willing to pay 7.8% of wages for a WFH option, which is almost exactly what they find in a survey DCE with a hypothetical scenario (8.4%) and in-line with what Lewandowski, Lipowska, and Smoter (2022) find in a DCE on a sample of workers in Poland for the option to work from home 2-3 days per week (7.3%).⁴⁰ In another field experiment where random choices of wages and hours were embedded in a real application process, Mas and Pallais (2019) estimate average value of non-work time at 60% of the wage for unemployed applicants. In their accompanying survey the same hypothetical choices they estimated an average value of non-work time of 61% of the wage. Another way to validate DCEs in surveys is to compare answers to actual choices made by participants. Wiswall and Zafar (2018), Maestas et al. (2023), and Mas and Pallais (2017) also show that survey responses align with observed choices of respondents. For example, Wiswall and Zafar (2018) elicit preferences of undergraduates over work flexibility and job stability and they find that respondent choices made on the survey relate to college major choices and to actual job choices reported in a follow-up survey four years after graduation.

Why do implied valuations in DCEs in the work context align so well when in other settings, such as environmental economics, they are checkered with inconsistencies, as reported by the influential survey of Diamond and Hausman (1994)? One clue is in Diamond and Hausman (1994), who write that they believe "that the internal consistency problems come from an absence of preferences, not a flaw in survey methodology." While respondents tend to have an absence of preferences over many questions on the environment ("How much do you value the bog turtle?"), people have clearly defined preferences and opinions about work. The types of questions DCEs ask about the workplace mimic the types of decisions people make throughout their careers. In my experience, based on feedback on surveys I have administered, people do not have

 $^{^{40}}$ Traditional survey questions where respondents self-report find the same magnitude response, such as Barrero, Bloom, and Davis (2021) (8.2%) and De Fraja et al. (2024) (8%).

difficulty with these questions, and tend to enjoy answering them.

Another indicator of the validity of information gleaned from DCEs is their use in the market. Companies, with real stakes, are now using them to design benefit plans and schedules (DeBellis 2018). An interesting question is whether the increasing use has resulted in employers extracting surplus from workers. Without surveys, employers could only guess how employees valued benefit packages, but accurate WTP measures allow employers to calibrate their plans to leave workers at the point of indifference. Freeman (1981) conjectured that employees prefer not to give information to employers for this reason, leaving unions to better represent preferences. This raises questions about incentive compatibility in employer surveys, something that does not appear to present a major problem in researcher-administered surveys where respondents do not have an incentive to conceal their preferences.⁴¹

The labor literature using DCEs has now generated a wealth of information on the workforce including on teachers (Johnston 2020; Tsao 2025), shift premia (Desiere and Walter 2023), autonomy and teamwork (Non et al. 2022), supervisor respect (Dube, Naidu, and Reich 2022), and hostile climates (Collis and Van Effenterre 2025), just to name a few. Maestas et al. (2023) has been particularly useful because their DCE is in a representative (US) sample and asks about a wide-range of non-wage amenities and working conditions.

Having this view into worker preferences opens up many possibilities as it offers potential to incorporate quantified preferences in models to learn about other features of the economy. For example, the estimates from these DCEs could be used as target moments in macro-models. The simple exercise I outlined in the context of Mas and Pallais (2017) on estimating employer WFH costs using information on WTP and observed quantities is an example of how this can be done. Drake, Thakral, and Tô (2022) conduct a more systematic exercise using estimated preferences to learn about the wage structure. Colonnelli et al. (2023) use estimates on preferences for ESG practices from a survey experiment to structurally estimate the joint distribution of TFP, demographic productivity multipliers, and nonwage amenities at the firm level.

8.2 MWTP Estimation with Non-Experimental Offer Data

Mas and Pallais (2017) is an ideal case where experimental variation in the wage allows not just for the estimation of MWTP, but of the full amenity demand function. In this

⁴¹There are methodological tools that can be used when truthful revelation of preferences is a concern. An illustrative application is Colonnelli et al. (2023) who embed the preference-elicitation process for employer ESG signals within a major job-search website, tying participants' stated choices to real-time job recommendations and hence rewarding truthful reporting.

section, I discuss non-experimental approaches.

Let a firm or job consist of an offered wage w_{ij} and its attributes a_{jk} . Workers have utility over these features, and there is also an idiosyncratic component of utility specific to the worker and firm match. Suppose we have data on a set of workers, and observe all of their offers, both accepted and not accepted. We also observe the characteristics of their alternatives. Such a dataset is quite special, and opens up possibilities to learn about preferences. However, using data that are equilibrium objects as opposed to varying experimentally presents some challenges to identification.

In this setting each worker i chooses between different jobs (or firms) according to a logit-type random-utility model. In such a model, the utility of worker i at firm j is

$$U_{ij} = \beta \ln(w_{ij}) + \sum_{k=1}^{K} \gamma_k a_{jk} + \varepsilon_{ij}.$$
(9)

 ε_{ij} is an idiosyncratic shock, which is often assumed to be drawn from an i.i.d. extremevalue type-I (EV1) distribution, and we will proceed this way. Since wages appear in logs, the ratio of the γ parameters to β give the MWTP as a share of the wage: $\frac{MWTP_k}{w_{ij}} = \frac{\gamma_k}{\beta}$.

The offer wage varies with individual, firm and match characteristics. Assuming these are separable we can write:

$$ln(w_{ij}) = \alpha_i + \theta_j + \mu_{ij}.$$

The component θ_j comes from the firm's cost-minimization problem, as in the benchmark model, where they set wages and amenities to match a target utility level \overline{V}_j . The term μ_{ij} is a match component and α_i reflects the portable individual wage component that captures skill. We have evidence from Hall and Mueller (2018) that the latter two components have explanatory power. They analyze offer data for unemployment insurance recipients and estimate that 25% of the variation in wage offers is due to differences in offers for workers with the same productivity.

This model can in principle be estimated with a variant of a logit (e.g., binary logit, multinomial logit or mixed logit depending on the exact structure of the choice and regressors) using data on job offers.⁴² But there are a number of qualifications that are related to the usual concerns when estimating a Berry, Levinsohn, and Pakes (1995) (BLP) model, but in the labor context.

The first assumption is that match-specific utility ε_{ij} has to be independent of the match term in wages μ_{ij} . This assumption holds tautologically under wage posting,

⁴²As in Berry, Levinsohn, and Pakes (1995), the model can be estimated in two-steps, first by estimating firm-specific effects and then regressing the firm effects on firm characteristics.

since there is no match component in the wage, or if the bargained wage is not a function of a worker's idiosyncratic utility.

The second assumption is that all firm characteristics that are valued by workers are included in the set of amenities explicitly modeled. To the extent that there are missing characteristics, the γ_k parameters are interpretable as preferences over bundles of characteristics that are correlated to the relevant $\alpha_j k$. The cost-minimization problem also implies that unobserved characteristics will load onto the wage both because better unobserved amenities \implies higher target utility $\overline{V}_j \implies$ higher firm component of wage (since firms vary wages and amenities jointly to match target utility) and higher average skill.

Given these strong assumptions, a valid instrument for the wage can be helpful, though often difficult to find. It would shift labor demand without also shifting workers' unobserved preferences. This instrument would be at the firm level, similar to an IV for estimating the labor supply elasticity to the firm—for example, changes in union contracts, firm-specific demand shocks, heterogeneous effects in the tax code, or exogeneous changes in input prices that pass through to wages via changes in profitability and rent-sharing. As in BLP, we could contemplate using Hausman (1978)-style instruments if we had data on the same firm operating in multiple markets and use the wage in city B as an instrument in wage in city A with the idea that the driver of these correlated wages are a firm's cost structure rather than market-specific labor supply.⁴³ In practice, much care has to be taken to ensure that the chosen instrument is not also affecting the other side of the market. For example, an infrastructure upgrade like improved roads would affect both labor demand (lower distribution costs) and supply (easier commute).

Absent an external instrument, if we are willing to maintain the first assumption that $\text{Cov}(\mu_{ij}, \varepsilon_{ij}) = 0$, and we have extensive data on offers to the same individuals, we can estimate the match component μ_{ij} and use that as an instrument for the wage. Using a match effect as an IV ensures that the variation used to estimate the relationship between wage and acceptance probability is orthogonal to all other firm characteristics. Intuitively, the reduced form of this instrument can be estimated to capture the relationship of the wage on job acceptance while including firm fixed effects that control for all stable features of the firm, including amenities possibly correlated to the wage. Of course, the assumption that the match term on wages is uncorrelated to match-specific utility is strong.

A final issue for consideration is non-random offers. A dataset with offer data will excludes all firms where a worker did not apply, and firms that did not make an offer

⁴³The first-stage of this instrument should be strong given the findings in Hazell et al. (2022) on national wage setting.

to the worker. This is a sample selection problem and in principle could be addressed with a control function approach assuming that we have some Z variable that predicts a job offer and can be excluded from the second stage logit.

All of this discussion assumes that we even have data on offers. In fact, these are quite rare. Roussille and Scuderi (2023) use offer data from a U.S job search platform to differentiate between different models of market power.⁴⁴ A promising approach absent offer data is to use either survey or job flow data, as in Caldwell, Haegele, and Heining (2024) to construct consideration sets for workers—the set of feasible or actually considered options. These sets address part of the data requirement for the choice problem. Then some additional structure or information would be required on the firm side to model the probability that an offer is issued. Caldwell, Haegele, and Heining (2025) side-step this problem by conducting a discrete choice experiment over firm choices from worker-defined consideration sets to estimate the value of a firm.

A related problem are occupational choice models, where the amenities are defined at the occupational level. This type of model was first estimated by Boskin (1974), who used occupational choice to estimate relative weights for discounted present value of potential future earnings, retraining costs for occupational entry, and present value of expected earnings foregone due to unemployment. Conditional on the first, the latter two are hypothesized to have more weight when there are imperfect capital markets. A second example is DeLeire and Levy (2004), who employ a choice model to estimate the relationship between fatality risk and occupational choice by family structure.

On the surface, occupational choice models have fewer data requirements than choices over offers, since occupation is recorded in most labor force surveys. Unfortunately, implicit in these models is an untenable assumption that all occupations have a positive probability of selection. This assumption is unrealistic once we consider credential and qualification requirements, and self-selection as some individuals might not even consider occupations deemed unsuitable (e.g., due to physical demands). Knowing an individual's consideration set is crucial. If we ignore the fact that certain alternatives are not in a worker's consideration set, we run the risk of misspecifying the probability of choosing an observed option, biasing estimated preference parameters. A solution is information on the consideration set is not available are "consideration instruments" that affect the the probability that an occupation is in the choice set, but are excluded from worker utility. These would be used in a control function to correct for selection. See Dubé, Hortaçsu, and Joo (2021) for an application of this approach in the marketing

⁴⁴Stern (2004) has offer data but uses it to estimate a hedonic model rather than a choice model. It would be interesting to see in his setting how MWTP estimates from observed choices compares to his hedonic estimates discussed previously.

context.

Quasi-Experimental Approaches There are several papers that have estimated job choice problems using quasi-experimental variation. The common thread is that each worker is choosing between two or more options with varying relative wages and, in these cases, risk.

Greenberg et al. (2021) use Army reenlistment decisions to estimate the wage-risk premium, and thus the VSL. This approach is the revealed-preference analogue to standard hedonic methods; instead of relating wages to risk directly, it examines how individuals choose among alternatives differing in risk and wages. The Army setting illustrates why revealed preference can be advantageous. Reenlistment bonuses vary with attrition and occupation vacancies, so riskier occupations presumably command higher bonuses to attract soldiers. Yet these bonuses are determined administratively rather than by a competitive labor market. Administrators can also respond to vacancies in other ways such as more intensive recruiting, investing in equipment, or changing operations to deprioritize hard to fill occupations. Bonuses might be set high for non-combat roles that are high priority but not necessarily high risk. It is unclear what a hedonic model would capture. By contrast, a revealed-preference approach infers risk valuation from actual choices, regardless of whether wages or bonuses are set competitively.

In their benchmark specification, Greenberg et al. estimate a logit model of a soldier's reenlistment decision as a function of the reenlistment bonus and the fatality risk associated with their military occupation. The data span 2002–2011, a period marked by elevated casualties—particularly in combat operations—due to U.S. military involvement in Afghanistan and Iraq. Their results show a strong, positive relationship between reenlistment and the bonus, and a strong negative relationship between reenlistment and fatality risk, yielding VSL estimates in the \$700,000–\$1 million range, substantially lower than those in earlier hedonic wage-risk studies.

Another paper from the field that is in spirit a choice problem is Anelli and Koenig (2021). They adopt a revealed-preference framework by analyzing frontline hourly workers' decision to keep working or reduce earnings under varying COVID-19 risk levels. They use discontinuities in Federal Pandemic Unemployment Compensation (FPUC) benefits, which create a choice margin: qualify for higher UI benefits by cutting hours versus continue working at higher risk. Because COVID-19 risk functions like an implicit tax on work, observing how workers bunch below the earnings threshold—especially in higher-risk environments—reveals how much they value safety. From this behavior, Anelli and Koenig back out a WTP for risk reduction and derive a corresponding value of a statistical life of \$6.9 million.

8.3 Using Worker Mobility to estimate MWTP for Job Characteristics

Most datasets do not have data on multiple offers, but Gronberg and Reed (1994) showed that MWTP can be estimated with only data on voluntary worker separations by assuming a simple search model. Jobs are described by a vector of characteristics $X = (w, a_1, ..., a_n)$, which includes the wage. In a Mortensen (1986) search environment, exogeneous job destruction is given by δ , λ is the job offer arrival rate, $F(\cdot)$ is the cdf of utility of job offers, and *b* is reservation utility. Let V(X) be a quasi-concave utility function denoting the utility flow a worker receives from a job with characteristics *X*. $s^*(V(X))$ determines search intensity as a function of job value. The hazard rate is

$$h(V(X)) = \delta + \lambda s^*(V(X)) \left[1 - F(V(X)) \right], \text{ for all } V(X) \ge b.$$

The first piece of this expression is due to the worker having to leave their job because of an exogenous job destruction event. The second comes from the worker being able to receive a job offer (which happens with probability λ), which depends on search intensity s(V(X)) and provided that the new job has higher utility than the current one which occurs with probability [1 - F(V(X))].

We can examine how a small change in a particular job attribute a_k shifts the hazard rate. Using $V \equiv V(X)$ for brevity, formally,

$$\frac{\partial h}{\partial a_k} = \frac{\partial V}{\partial a_k} \left[\frac{\partial s^*}{\partial V} \lambda \left(1 - F(V) \right) + \frac{\partial \left(1 - F(V) \right)}{\partial V} \lambda s^*(V) \right], \quad k = 1, \dots, n$$

Note that in the above expression, the term in brackets does not depend on any of the characteristics, or the wage, other than through utility. Therefore, if we take the ratio of the partial derivatives of the hazards for an amenity and the wage, that term drops out and we isolate the partial of utility. Taking ratios then gives the MWTP:

$$\text{MWTP}_{k}(X) = \frac{\partial h(X) / \partial a_{k}}{\partial h(X) / \partial w} = \frac{\partial V / \partial a_{k}}{\partial V / \partial w}$$

As we have seen, unobserved firm characteristics are a (the!) major confounding factor in hedonic wage regressions. This approach delivers unbiased MWTP even with these unobservables. Hwang, Mortensen, and Reed (1998) show that unlike hedonic models, this approach is robust to frictions.

There are several assumptions that deliver this result. The first is that search intensity is not a function of the wage or job characteristics other than through utility. Van Ommeren, Van Den Berg, and Gorter (2000) show that if search is endogenous, that is, a worker can optimize on it, then this conditon does not hold.

A second assumption is that that exogenous job destruction rate does not vary with the wage and job characteristic. This would happen if business cycles have differential impacts on workers with different amenities, or if there are efficiency wages such that when workers have higher utility they are less likely to shirk.

A third assumption is that the characteristic X does not affect the hazard except through utility. Working from home may present a violation of this assumption if being able to work from home is both beneficial in utility and makes it easier for workers to look for other jobs. Bonhomme, Jolivet, and Leuven (2016) argue this assumption is potentially violated in the context of Dutch primary school teachers because amenities affect job utility and can constrain job opportunities, for example if job offers are a function of the socioeconomic level of the school.

The Gronberg and Reed (1994) approach has been used in a number of contexts. For instance, Van Ommeren, Van Den Berg, and Gorter (2000) used it with job duration data to estimate how workers value commute time. They estimate a duration model as a function of wage and commute; the ratio of these two coefficients implies that workers value their commute time at 50 percent of their wage. That is, a worker would have to be paid half their hourly wage for each additional hour of commute. This estimate is substantially higher than the MWTP from a hedonic model, a pattern that is generally found, giving credence to the Hwang, Mortensen, and Reed (1998) model. In a survey of Walmart workers, Dube, Naidu, and Reich (2022) ask about hypothetical quit decisions when showing them characteristics of other jobs. The resulting MWTP are effectively Gronberg and Reed (1994) estimates.

Using Information on Destination and Receiving Jobs In some household surveys we can observe origin and destination information for job movers. These differences in characteristics between firms give additional information beyond the separation probability. It may be tempting to compare changes in the wage to changes in job characteristics to estimate a MWTP, however, this approach is subject to selection bias. The sample is selected on movers for whom the wage change was sufficiently large to compensate them for the utility loss due to more disamenities or, conversely, the new amenities were sufficiently attractive to motivate a move even for a large wage reduction. The problem, ultimately, is that we do not observe wage and amenity offers for stayers.⁴⁵

Bonhomme and Jolivet (2009) impose structure to address the selection problem.

⁴⁵Villanueva (2007) shows that even with the selection bias, the MWTP can still be bounded. However, the bounds in his approach tend to be large.

They estimate workers' WTP for non-wage job characteristics, such as amenities, using observed voluntary job transitions. They only observe accepted job offers rather than the full distribution of offers workers receive. Selection bias arises as observed offers are conditional on workers' decision rules. Their solution is to use an estimate of the offer distribution to infer the probability of transitioning. The decision to switch jobs is modeled using a latent utility framework. The worker accepts an offer if the utility from the new job exceeds the utility of staying, net of transition costs. This decision rule can be expressed as:

$$\Phi^{-1}\big(P(z_t=1|w^*,a^*)\big)=\frac{1}{\sigma_z}\big(w^*+\beta a^*-\tau(x,\theta)\big),$$

where $z_t = 1$ indicates a voluntary job transition, w^* and a^* are the wage and amenity offers, β captures the worker's WTP for the amenity, and $\tau(x, \theta)$ summarizes the effects of individual characteristics and unobserved heterogeneity. This equation connects the observed probability of transitioning to the worker's trade-off between wages and amenities.

To address the selection problem, the model leverages the relationship between observed transitions and the underlying distribution of offers. Specifically, the probability of a transition is adjusted using Bayes' rule:

$$P(z_t = 1 | w^*, a^*) = \frac{\ell_z(w^*, a^* | z_t = 1)}{\ell^*(w^*, a^*)} P(z_t = 1),$$

where ℓ_z represents the density of accepted offers (conditional on transitioning), and ℓ^* represents the overall distribution of offers in the labor market. By estimating ℓ^* , the model corrects for the selection bias inherent in observing only accepted offers. The offer distribution is inferred using constrained transitions (e.g., layoffs or moves due to firm closures), which are assumed to provide an unbiased sample of market offers. Once the selection-corrected expression for the transition probability is established, δ^* can be recovered as the coefficient on amenities a^* in the estimated decision rule. This allows them to directly quantify workers' WTP for non-wage job characteristics. As in papers that use the Gronberg and Reed (1994) approach, Bonhomme and Jolivet (2009) find substantially larger MWTP values in absolute value relative to hedonic regressions. This again would be because choice approaches are robust to frictions.

8.4 Using Flows to Estimate Firm Value

If we have employer-employee matched data, we can not only observe a separation, but track where the worker moves. These flows between firms are informative about the value of the firm. While the previous approaches are oriented towards valuing specific features of firms, with employer-employee matched data we can attempt to construct relative total firm values (wages+amenities). The key reference here is Sorkin (2018). With a measure of firm value in hand, we can in a second step use this measure to estimates the value of characteristics, or investigate how the distribution of firm values contributes to, say, inequality.

Suppose that a firm or job can be defined by a component of common value. This value represents all pecuniary and non-pecuniary parts of the employment relationship so that utility is defined as:

$$U_{ij} = V_j + \varepsilon_{ij}.$$

As before ε_{ij} is an idiosyncratic shock drawn from an i.i.d. EV1 distribution, and V_j is the utility of firm *j* that is common to all workers.

By the logit structure, the probability that any worker i chooses firm k over firm j is

$$P(\text{choose } k \text{ over } j) = \frac{\exp(V_k)}{\exp(V_k) + \exp(V_j)} = \frac{M_{kj}}{N}$$

When we observe, in aggregate, that M_{kj} workers choose k instead of j (out of N who actively compare these two firms). Abstracting from firm size, the relative number of people that chose k over j (M_{kj}) versus j over k (M_{jk}) is the measure of the relative common value of the firms:

$$V_k - V_j = \ln\left(\frac{M_{kj}}{M_{jk}}\right),\tag{10}$$

where M_{jk} denotes the count of those who choose j instead of k.⁴⁶

If we repeat this logic across many pairs (k, j) we can in principle recover each V_k relative to an outside option on a common utility scale by maximum likelihood using the Bradley-Terry-Luce (BTL) model (Bradley and Terry 1952; Luce 1959). In the context of the labor market, however, this approach is generally infeasible because of the sparsity of the transition matrix. The bilateral flows will have many structural zeros, complicating estimation. An additional problem is that there is also a decision rule on the firm side as to whether to make an offer.

Sorkin (2018)'s solution collapses all the pairwise conditions into one equation per

⁴⁶Voluntary transitions of workers between employers are required. Papers that use these flows with matched employer-employee data typically approximate voluntary moves by making sure that there are no quarters with zero earnings gaps after a job move.

firm, massively reducing dimensionality. Concretely, expression (10) implies that:

$$M_{kj} \exp(V_j) = M_{jk} \exp(V_k), \quad \forall j.$$
$$\implies \exp(V_k) = \frac{\sum_{j \in \mathscr{E}} M_{kj} \exp(V_j)}{\sum_{j \in \mathscr{E}} M_{jk}}, \tag{11}$$

where \mathscr{E} is the set of all firms. In this expression each firm's (exponentiated) value equals its value-weighted inflow divided by its total outflow. By aggregating over all sources *j* into a single inflow term, one no longer needs direct observations of every (j,k) transition, thereby avoiding the problem of zero flows in sparse data. This system will have a unique solution under relatively mild connectivity conditions, and can be efficiently solved numerically even when the number of firms is very large.

Sorkin (2018)'s approach is motivated by the "PageRank" algorithm (Page 1999), because a high-value firm draw workers from other high-value firms and loses few workers, similar in spirit to the way the PageRank algorithm ranks websites. However, Sorkin's approach is much more akin to another spectral ranking algorithm, Negahban, Oh, and Shah (2017)'s RankCentrality. RankCentrality is an iterative rank aggregation algorithm for discovering strength of items when the data adhere to the pairwise comparisons of the BTL model. RankCentrality is defined as the stationary distribution of a random walk on a weighted directed comparisons graph. This stationary distribution is precisely the fixed point to equation (11), where the M_{kj} terms are the elements of the comparisons graph. There is a stationary distribution when the elements are in a strongly connected set, as defined by Sorkin.

While this is an elegant solution for estimating the latent utility of firms, the model of the labor market that leads to this result is highly simplified. The value M_{kj} represents the number of workers who prefer k to j among all workers who are potential movers. In practice, with standard administrative datasets, we do not observe this value. Instead we observe the number of workers who prefer k to j conditional on being employed at j and receive an offer from k. We do not observe the number of workers at j who prefer k but do not receive an offer, or the number of workers at j who did receive an offer but preferred j to k. This presents a problem because firms, especially high value firms, may make limited offers, and have a queue of workers who would join the firm, if offered. Without taking into account differential offer rates, these firms would appear less attractive than they are.

This problem can be addressed by imposing structure from a Burdett and Mortensen

(1998) model. In steady state:

$$M_{jk} = \underbrace{g_k W(1 - \delta_k)(1 - \rho_k)}_{\text{# workers staying at }k} \times \underbrace{\lambda_1 f_j}_{\text{prob. of receiving an offer from }j} \times \underbrace{\frac{\exp(V_e(v_j))}_{\exp(V_e(v_j)) + \exp(V_e(v_k))}}_{\text{prob. of accepting }j}$$

where g_k is the share of all workers employed at k, W is the total number of workers, δ_k and ρ_k are exogeneous separation rates, and $\lambda_1 f_j$ is the probability an employed worker at k gets an offer from j. We get a similar expression from $j \rightarrow k$, and taking the ratio of the two yields

$$\frac{M_{jk}}{M_{kj}} = \underbrace{\frac{f_j}{f_k}}_{\text{relative offers}} \times \underbrace{\frac{g_k(1-\delta_k)(1-\rho_k)}{g_j(1-\delta_j)(1-\rho_j)}}_{\text{effective sizes}} \times \underbrace{\frac{\exp(V_e(v_j))}{\exp(V_e(v_k))}}_{\text{relative values}}.$$

Sorkin defines the "flow-relevant" utility as

$$\exp(\tilde{V}_j) \equiv \frac{f_j \exp(V_j)}{g_j (1 - \delta_j) (1 - \rho_j)}$$

It follows that

$$\frac{M_{jk}}{M_{kj}} = \frac{\exp(\tilde{V}_j)}{\exp(\tilde{V}_k)},$$

and we can apply this expression to the fixed-point problem in equation (11), retrieving $\exp(\tilde{V}_k)$. The firm value, $\exp(V_k)$, is then obtained by multiplying the resulting quantity by $\frac{g_j}{f_i}$,

$$\widehat{V}_j = \frac{g_j}{f_j} \exp(\widetilde{V}_k).$$

Here, g_j represents firm size and is readily available in employer-employee matched data, while the offer rate f_j needs to be estimated. As in Bonhomme and Jolivet (2009), offer rates are computed using constrained transitions. In Sorkin's case, these are the relative rates of hires out of non-employment. Since individuals in non-employment have limited outside options, observing that a firm hires more from non-employment suggests that it is making more offers. There is an assumption, then, that the offer rate to employed workers is the same. An additional necessary assumption is that everyone has an equal probability of getting an offer from a given firm. We have to rule out scenarios where, say, low value firms only make offers to job seekers who they are confident will accept and appear to have high value because of a resulting high acceptance rate.⁴⁷

⁴⁷An alternative measure of firm value is from Bagger and Lentz (2019), which says a higher ranked firm hires a lower share out of non-employment and a higher share of workers from other firms.
Another limitation of the approach is that everyone has essentially the same ranking of firms up to an EV1 error. Roussille and Scuderi (2023) relax that assumption by allowing for three distinct "job ladders" and find evidence to support this heterogeneity.

Though this RankCentrality measure of firm value does not use wages in any way, Sorkin (2018), Lachowska et al. (2023b), and Lagos (2024) all find that it is strongly related to a firm's wage component in an AKM model, as would be expected if it measured job value. The main substantive findings in Sorkin is that the measure of value varies a lot across firms. While wage variation may just reflect compensating differentials and thus not violate a law of one price, the observation of significant utility variation is decisive on that point.

8.4.1 Application to Collective Bargaining Agreements

Lagos (2024) uses this RankCentrality measure to quantify wage-equivalent values of collective bargaining agreement clauses. He estimates $\Delta \hat{V}_j = \beta_1 \Delta \hat{\psi}_j + \sum_{k=1}^K \beta_k \Delta a_{kj} + \varepsilon_j$, where the a_{kj} is contract clause k at firm j and $\hat{\psi}_j$ is the change in the firm-component of wage defined over non-overlapping five year intervals. Using a measure of firm utility as a dependent variable requires fewer assumptions than the Gronberg and Reed (1994) job separation approach because in the latter we have to assume that job characteristics only affect separations through utility, that is, they cannot have an independent affect on mobility. If we have a measure of firm utility, job characteristics can independently affect separation probabilities and still yield consistent estimates.⁴⁸

Lagos' findings include: Workers value restricting the duration of probationary periods and temporary contracts at 2.2% of wages; advance payment of year-end bonuses for vacations valued at 3.3% of wages; and paid leave for deaths, blood donations, marriage, and paternity valued at 7.8% of wages. He finds negative WTP estimates for other clauses, such as a requirement that accidents involving leave must be reported to the union within 24 hours (-6.4% WTP) and a requirement that employers inform workers about vacation periods at least 30 days in advance (-5.1% WTP).

These estimates are interesting in light of Freeman (1981), who argues that fringe benefit allocations at unionized employers are a good indicator of worker preferences because accurate information emerges from the collective bargaining process. The results in Lagos (2024) suggest limits to that interpretation because there are some provisions that are not valued. The pattern appears that workers value contract clauses

⁴⁸In principle, we can use any measure of firm utility as the dependent variable, like average job satisfaction or other measures of subjective well-being. However, Benjamin et al. (2014) show that choice-based approaches, like Sorkin's, result in more accurate MRS estimates than subjective measures of utility.

that have direct benefits, and do not value provisions whose sole purpose are potentially costly administrative obligations for the employer. The finding is reminiscent of Boudreau (2024), who found negative work satisfaction responses to the implementation of safety committees. One interpretation of these findings is that the incentives of union officers are not completely aligned with the rank-and-file as officers put excessive weight on administrative procedures.

8.4.2 Application to Estimation of Work Hour Constraints

Lachowska et al. (2023b) estimate a model of the form

$$\widehat{V}_j = c + \theta \ln(\text{wage}_j) + \beta \ln(h_j) + \varepsilon_j, \qquad (12)$$

where \hat{V}_j is a RankCentrality measure, and h_j is a measure of firm work hours that varies due to policy, as opposed to worker preferences. They show that the ratio of the MRS between earnings and work hours to the wage is: $\frac{MRS_{e,h}}{w} = -\frac{\beta-\theta}{\theta}$. If work hours are unconstrained, as in the neoclassical models of labor supply, this ratio should equal 1. The authors show that it is well below 1 on average, implying that there are equilibrium work hour constraints.

Equilibrium constraints in work hours emerge naturally out of any model where firms have preferences over average hours worked due, for example, to scheduling constraints or fatigue. Figure 5 shows how this happens. Worker indifference curves in wage and hour space are u-shaped. If hours are too low people would take a pay-cut to work more. If hours are too high they would take a pay cut to work less. If firms have preferences over average hours, they will require either a higher or lower wage (depending on their technology) to increase hours to maintain profits. In the figure, the isoprofit for the firm is drawn downward sloping implying that extra hours are costly to the firm. In that case, they will offer constrained hours at point a, with wage of w_1 . At this wage, unconstrained hours are given by the tangency point b at the minimum of indifference curve corresponding to utility V_2 . Using the RankCentrality measure, Lachowska et al. estimate unconstrained hours empirically by the maximum average RankCentrality value for a fixed wage bin. Using this value, and the estimate of θ in equation (12), they can determine the wage premium required to achieve the same utility at unconstrained hours but from the constrained hours. This premium is shown as $w_2 = w_1 \exp\left(\frac{V_2 - V_1}{\theta}\right)$ in the figure. They find that, on average, the wage premium required to equalize utility for constrained hours to unconstrained hours is 12%.





8.5 Equilibrium Models

The revealed preference approach has been embedded in structural models that seek to explain the wage structure. Amenities in these models typically play the role of the residual, explaining anything that is not rationalized by observed firm, worker or market characteristics given a specified model.⁴⁹

Sullivan and To (2014), Sorkin (2018), Taber and Vejlin (2020), and Morchio and Moser (2024) rely on information on job transitions in frictional model to quantify the value of non-wage amenities, with the intuition that a move to a lower paying job conveys information on non-pecuniary aspects of the job. Hall and Mueller (2018) also model a frictional market but instead use offer data for unemployment insurance recipients. They infer the importance of non-wage amenities using the share of job offers accepted below a self-reported reservation wage. Amenities can rationalize, for example, why job acceptance is kinked but not discontinuous at the self-reported reservation wage.⁵⁰

By contrast, Lamadon, Mogstad, and Setzler (2022) do not rely on search frictions

⁴⁹Reassuringly, when data on actual amenities are available, these quantities tend to be correlated to be correlated to these residual elements in expected ways, as shown by Morchio and Moser (2024).

⁵⁰The importance of non-wage amenities, and the fact that they are correlated to the wage, suggests that, in general, we have to be careful about how to interpret self-reported reservation wage measures. What are respondents assuming about all other aspects of work when providing a wage?

and instead use a horizontal amenity assumption and resulting match specific utility in a choice framework to explain how workers allocate to firms. In their model, workers in a given skill group can choose any firm and choose the one that gives highest utility given idiosyncratic preferences, wages and vertical amenities. They also have data on firm productivity from which they can estimate rent-sharing which arises due to variation in markdowns. More productive firms and firms endowed with better vertical amenities attract more workers. Unlike in search models where amenities affect job-to-job transitions, here the relevant moment that pins down vertical amenities is scale. Vertical amenities explain why a firm scales beyond what can be explained by productivity and rent-sharing.

One of the principle goals of equilibrium models is to quantify how non-wage amenities explain variation in overall utility or wages. In an equilibrium model this can be achieved by turning off one of the mechanisms in the model, as in Taber and Vejlin (2020), by conducting variance decompositions on different components of the wage based on estimated model parameters, as in Lamadon, Mogstad, and Setzler (2022), or by directly estimating the variances of different wage components and overall job value given target moments, as in Hall and Mueller (2018). Taber and Vejlin (2020) shut off compensating differentials by having all moves be based on wage rather than non-pecuniary elements and then compare the variance of the wage and utility to the observed variances. Using data from Denmark, they estimate that non-wage amenities explain more than half of the overall variance in flow utility. Using data from the State of New Jersey, Hall and Mueller (2018) conclude that total job value dispersion is 50% higher than wage dispersion. Using data from Brazil, Monaco and Pierce (2015) estimate that the amenity share in total compensation is approximately 50%.

It is informative to compare decompositions on the overall value of jobs with estimates from papers that have directly estimated non-wage amenities and quantified their value. Maestas et al. (2023) find that the 90-10 log compensation differential, inclusive of non-wage amenity values, is 11% larger than the 90-10 wage differential. Their amenities do not include fringe benefits, but Pierce (2001) finds a similar magnitude contribution from benefits.⁵¹ Maestas et al. (2023) also find that mean amenity value ranges from 35% to 45% of job value excluding fringe benefits, with higher share at the top of the income distribution. Fringe benefits are approximately 30% of wages, or about 17% of total compensation. Wages therefore, are about 55% of total compensation.

A related question that can be addressed in equilibrium models is whether varia-

⁵¹The Pierce (2001) estimate is based on benefit average costs, which should be lower than how much they are valued.

tion in wages across firms is explained by compensating variation due to differences in amenity provision. If so, then AKM firm effects or inter-industry wage differentials may be fully compatible with a competitive model. Krueger and Summers (1988) was an early paper to test this mechanism, comparing the variance of industry wage differentials with and without working condition controls. They found no support that interindustry wage differentials were explained by differences in non-pecuniary benefits. In the Taber and Vejlin (2020) model, when moves are based soley on wages, the variance of wages is reduced by 16% relative to when moves are also based on amenities. Sorkin (2018) decomposes AKM variation in wages due to compensating differentials by attributing the residual variance from a regression of the firm wage on the RankCentrality measure of firm value. The idea is that holding value fixed, any variation in the wage is offsetting amenities. Concretely, assume that the true data-generating process (DGP) for firm wages is $w_j = \beta V_j + \varepsilon_j$, where total firm utility is defined as $V_j = w_j + \gamma a_j$. Matching terms from this identity implies $\beta = 1$ and $\varepsilon_i = -\gamma a_i$. Sorkin estimates this relationship in the data and interprets the estimated residual variance as the contribution of non-wage amenities to overall wage dispersion. This exercise corresponds to the thought experiment of how much wages would vary if amenities were equalized at a given target utility level, in which case compensating differentials can only increase wage dispersion. If instead amenities are equalized across all firms, the change in wage variance becomes $\operatorname{Var}(\varepsilon_i) - 2\gamma \operatorname{Cov}(V_i, a_i)$. Since search and monopsony models such as Hwang, Mortensen, and Reed (1998) predict a positive covariance of firm value and amenities, this second term is positive. Therefore, depending on the relative sizes of the residual variance and the covariance term, compensating differentials can contribute either positively or negatively to the overall variance of wages.

9 Conclusion

The recent literature on non-wage amenities summarized in this chapter points to several main conclusions:

- Models that incorporate imperfect competition are necessary to evaluate how workers value amenities, and how amenities affect the wage structure.
- Omitted measures of job and firm productivity are a primary source of bias in compensating differential estimates when there is rent sharing or other frictions.
- Revealed-preference approaches tend to be more robust to deviations of perfect competition when estimating preferences over amenities.

• Survey-based DCEs are a viable way to elicit preferences over workplace characteristics over which workers have well-articulated preferences.

Given the recent literature, these points may be relatively uncontroversial. However, these conclusions are a major departure from the conventional wisdom at the time of Rosen's (1986) seminal Handbook Chapter on equalizing differences where the focus was mostly on hedonics. While the relevance of imperfect competition is clear, that does not diminish the importance of the competitive benchmark. The logic of the competitive benchmark still operates in all of the models we reviewed. Agents still optimize, and market equilibria are determined, in part, by the primitives of the competitive model. However, without accounting for potential market power and frictions, the competitive model loses empirical relevance.

These conclusions are also compatible with compensating differentials. They clarify why non-wage amenities can have important implications for the wage structure even if naive tests fail to detect them. Though rare, designs that take seriously market-structure and confounders do find that jobs with more attractive attributes have offsetting wages.

References

- Abowd, John A and Thomas Lemieux (1993). "The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada". In: *The Quarterly Journal of Economics* 108.4, pp. 983–1014.
- Abowd, John M., Francis Kramarz, and David N. Margolis (Mar. 1999). "High Wage Workers and High Wage Firms". In: *Econometrica* 67.2, pp. 251–333.
- Agency for Healthcare Research and Quality (2025). *Mean expenditure per person, United States, 1996 to 2022. Medical Expenditure Panel Survey.* URL: https:// datatools.ahrq.gov/meps-hc/?tab=use-expenditures-and-population& dash=12.
- Akerlof, George A (1984). "Gift exchange and efficiency-wage theory: Four views". In: *The American Economic Review* 74.2, pp. 79–83.
- Aksoy, Cevat Giray, Jose Maria Barrero, Nicholas Bloom, Steven J. Davis, Mathias Dolls, and Pablo Zarate (Sept. 2022). "Working from Home Around the World". In: *NBER Working Paper Series*.
- Anelli, Massimo and Felix Koenig (Dec. 2021). "Willingness to pay for workplace safety". In: *CESifo Working Paper Series* 14919.
- Arnold, Daniel and Christopher M. Whaley (July 2020). "Who Pays for Health Care Costs? The Effects of Health Care Prices on Wages". In: *RAND Working Paper Series*.
- Arnold, Daniel R, Jaime S King, Brent D Fulton, Alexandra D Montague, Katherine L Gudiksen, Thomas L Greaney, and Richard M Scheffler (Feb. 2025). "New evidence on the impacts of cross-market hospital mergers on commercial prices and measures of quality". In: *Health Services Research* 60.1, e14291.
- Arold, Benjamin W, Elliott Ash, W Bentley MacLeod, and Suresh Naidu (2024). "Do words matter? the value of collective bargaining agreements". In: *Center for Law & Economics Working Paper Series* 6.
- Ashenfelter, Orley and Michael Greenstone (May 2004). "Estimating the value of a statistical life: The importance of omitted variables and publication bias". In: *American Economic Review* 94.2, pp. 454–460.
- Bagger, Jesper and Rasmus Lentz (Jan. 2019). "An empirical model of wage dispersion with sorting". In: *Review of Economic Studies* 86.1, pp. 153–190.
- Baicker, Katherine and Amitabh Chandra (July 2006). "The Labor Market Effects of Rising Health Insurance Premiums". In: *Journal of Labor Economics* 24.3, pp. 609– 634.

- Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis (2021). *Why working from home will stick.* Tech. rep. National Bureau of Economic Research.
- Benjamin, Daniel J, Ori Heffetz, Miles S Kimball, and Alex Rees-Jones (Nov. 2014).
 "Can marginal rates of substitution be inferred from happiness data? Evidence from residency choices". In: *American Economic Review* 104.11, pp. 3498–3528.
- Berry, Steven, James Levinsohn, and Ariel Pakes (July 1995). "Automobile Prices in Market Equilibrium". In: *Econometrica* 63.4, pp. 841–890.
- Black, Dan A, Jose Galdo, and Liqun Liu (June 2003). *How Robust are Hedonic Wage Estimates of the Price of Risk?* Final Report to the USEPA [R 829-43-001].
- Bloom, Nicholas, James Liang, John Roberts, and Zhichun Jenny Ying (Feb. 2015)."Does working from home work? Evidence from a Chinese experiment". In: *Quarterly Journal of Economics* 130.1, pp. 165–218.
- Bonhomme, Stéphane and Grégory Jolivet (June 2009). "The Pervasive Absence of Compensating Differentials". In: *Journal of Applied Econometrics* 24.5, pp. 763–795.
- Bonhomme, Stéphane, Grégory Jolivet, and Edwin Leuven (Aug. 2016). "School characteristics and teacher turnover: Assessing the role of preferences and opportunities". In: *The Economic Journal* 126.594, pp. 1342–1371.
- Boskin, Michael J (Mar. 1974). "A conditional logit model of occupational choice". In: *Journal of Political Economy* 82.2, Part 1, pp. 389–398.
- Boudreau, Laura (July 2024). "Multinational Enforcement of Labor Law: Experimental Evidence on Strengthening Occupational Safety and Health Committees". In: *Econometrica* 92.4, pp. 1269–1308.
- Bradley, Ralph Allan and Milton E Terry (Dec. 1952). "Rank analysis of incomplete block designs: The method of paired comparisons". In: *Biometrika* 39.3/4, pp. 324– 345.
- Brot-Goldberg, Zarek, Zack Cooper, Stuart V Craig, Lev R Klarnet, Ithai Lurie, and Corbin L Miller (2024). *Who pays for rising health care prices? evidence from hospital mergers*. Tech. rep. National Bureau of Economic Research.
- Brown, Charles (Feb. 1980). "Equalizing Differences in the Labor Market". In: *Quarterly Journal of Economics* 94.1, pp. 113–134.
- Burdett, Kenneth and Dale T. Mortensen (May 1998). "Wage Differentials, Employer Size, and Unemployment". In: *International Economic Review* 39.2, pp. 257–273.
- Bureau of Labor Statistics (Dec. 2024). Employer Costs for Employee Compensation September 2024. URL: https://www.bls.gov/news.release/pdf/ecec.pdf (visited on 12/17/2024).

- Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining (2024). "Firm Pay and Worker Search". In: UC Berkeley Department of Economics Working Paper.
- (2025). "Firm Pay, Amenities, and Inequality". In: UC Berkeley Department of Economics Working Paper.
- Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline (2018). "Firms and labor market inequality: Evidence and some theory". In: *Journal of Labor Economics* 36.S1, S13–S70.
- Case, Anne and Angus Deaton (2020). *Deaths of Despair and the Future of Capitalism*. Princeton, N.J.: Princeton University Press.
- Claxton, Gary, Matthew Rae, Aubrey Winger, Emma Wager, Ellen Byers, Jason Kerns, Greg Shmavonian, and Anthony Damico (Oct. 2024). *Employer Health Benefits:* 2024 Annual Survey. Tech. rep. San Francisco, CA: Kaiser Family Foundation.
- Clemens, Jeffrey (Winter 2021). "How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-employment Margins". In: *Journal of Economic Perspectives* 35.1, pp. 51–72.
- Collis, Manuela R and Clémentine Van Effenterre (2025). *Workplace Hostility*. Working Paper.
- Colonnelli, Emanuele, Timothy McQuade, Gabriel Ramos, Thomas Rauter, and Olivia Xiong (2023). *Polarizing Corporations: Does Talent Flow to*" *Good*" *Firms?* Tech. rep. National Bureau of Economic Research.
- Corradini, Viola, Lorenzo Lagos, and Garima Sharma (Sept. 2022). "Collective Bargaining for Women: How Unions Create Female-Friendly Jobs". In: *IZA Discussion Paper Series* 15552.
- Cropper, Maureen, Emily Joiner, and Alan Krupnick (Summer 2024). "Revisiting EPA's value per statistical life". In: *Review of Environmental Economics and Policy* 18.2, pp. 193–211.
- Dafny, Leemore, Kate Ho, and Robin S. Lee (Summer 2019). "The Price Effects of Cross-Market Mergers: Theory and Evidence from the Hospital Industry." In: *RAND Journal of Economics* 50.2, pp. 286–325.
- Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi (Aug. 2013). "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service". In: *Quarterly Journal of Economics* 128.3, pp. 1169–1218.
- De Fraja, Gianni, Jesse Matheson, Paul Mizen, James Rockey, Shivani Taneja, and Gregory Thwaites (2024). *Remote Work and Compensation Inequality*. Tech. rep.
- Deaton, Angus and John Muellbauer (June 1980). "An Almost Ideal Demand System". In: *American Economic Review* 70.3, pp. 312–326.

- DeBellis, Pete (2018). Surveying Employee Preferences for Rewards: A Primer. URL: https://www2.deloitte.com/content/dam/Deloitte/ca/Documents/ audit/ca-audit-abm-scotia-surveying-employee-preferences-forrewards.pdf.
- DeLeire, Thomas and Helen Levy (Oct. 2004). "Worker Sorting and the Risk of Death on the Job". In: *Journal of Labor Economics* 22.4, pp. 925–953.
- Desiere, Sam and Christian Walter (2023). "The shift premium: Evidence from a discrete choice experiment". In: *British Journal of Industrial Relations*.
- Diamond, Peter A and Jerry A Hausman (1994). "Contingent valuation: is some number better than no number?" In: *Journal of economic perspectives* 8.4, pp. 45–64.
- Drake, Marshall, Fernando Payró, Neil Thakral, and Linh T Tô (2024). "Bayesian Adaptive Choice Experiments".
- Drake, Marshall, Neil Thakral, and Linh T Tô (2022). *Wage differentials and the price of workplace flexibility*. Tech. rep. Mimeo.
- Dube, Arindrajit, Suresh Naidu, and Adam D. Reich (Sept. 2022). "Power and Dignity in the Low-Wage Labor Market: Evidence from Wal-Mart Workers". In: *NBER Working Paper Series* 30441.
- Dubé, Jean-Pierre, Ali Hortaçsu, and Joonhwi Joo (July 2021). "Random-coefficients logit demand estimation with zero-valued market shares". In: *Marketing Science* 40.4, pp. 637–660.
- Emanuel, Natalia and Emma Harrington (Sept. 2022). *Firm Frictions and The Payoffs* of Higher Pay: Labor Supply and Productivity Responses to a Voluntary Minimum Wage. Working Paper.
- (Oct. 2024). "Working Remotely? Selection, Treatment, and the Market for Remote Work". In: *American Economic Journal: Applied Economics* 16.4, pp. 528–559.
- Emanuel, Natalia, Emma Harrington, and Amanda Pallais (Nov. 2023). "The power of proximity to coworkers: training for tomorrow or productivity today?" In: *NBER Working Paper Series* 31880.
- Eriksson, Tor and Nicolai Kristensen (2014). "Wages or fringes? Some evidence on trade-offs and sorting". In: *Journal of Labor Economics* 32.4, pp. 899–928.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo FP Luttmer (Dec. 2019). "The value of medicaid: Interpreting results from the oregon health insurance experiment". In: *Journal of Political Economy* 127.6, pp. 2836–2874.
- Finkelstein, Amy, Casey McQuillan, Owen Zidar, and Eric Zwick (Spring 2023). "The health wedge and labor market inequality". In: *Brookings Papers on Economic Activity*, pp. 425–475.

- Freeman, Richard B (July 1981). "The effect of unionism on fringe benefits". In: *ILR Review* 34.4, pp. 489–509.
- Freeman, Richard B and Morris M Kleiner (May 2000). "Who benefits most from employee involvement: firms or workers?" In: *American Economic Review* 90.2, pp. 219–223.
- Goldin, Claudia and Lawrence F. Katz (Nov. 2011). "The Cost of Workplace Flexibility for High-Powered Professionals". In: *The Annals of the American Academy of Political and Social Science* 638, pp. 45–67.
- Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti (Jan. 2020). "Why Firms Offer Paid Parental Leave: An Exploratory Study". In: *NBER Working Paper Series*.
- Gray, George R, Donald W Myers, and Phyllis S Myers (May 1998). "Collective bargaining agreements: safety and health provisions". In: *Monthly Labor Review* 121.5, p. 13.
- Gray, Wayne B (1987). "The cost of regulation: OSHA, EPA and the productivity slowdown". In: *The American Economic Review* 77.5, pp. 998–1006.
- Greenberg, Kyle, Michael Greenstone, Stephen P Ryan, and Michael Yankovich (July 2021). "The heterogeneous value of a statistical life: Evidence from US Army reenlistment decisions". In: *NBER Working Paper Series* 29104.
- Greenstone, Michael (2017). "The continuing impact of Sherwin Rosen's "Hedonic prices and implicit markets: product differentiation in pure competition". In: *Journal of Political Economy* 125.6, pp. 1891–1902.
- Gronberg, Timothy J. and W. Robert Reed (Summer 1994). "Estimating Workers' Marginal Willingness to Pay for Job Attributes Using Duration Data". In: *Journal of Human Resources* 29.3, pp. 911–931.
- Gruber, Jonathan and Alan B. Krueger (1991). "The Incidence of Mandated Employer-Provided Insurance: Lessons from Workers' Compensation Insurance". In: *Tax Policy and the Economy* 5, pp. 111–143.
- Gruber, Jonathan and Michael Lettau (July 2004). "How elastic is the firm's demand for health insurance?" In: *Journal of Public Economics* 88.7-8, pp. 1273–1293.
- Guzman, Gloria and Melissa Kollar (Sept. 2023). *Income in the United States: 2022*. URL:https://www.census.gov/content/dam/Census/library/publications/ 2023/demo/p60-279.pdf.
- Hall, Robert E and Andreas I Mueller (Aug. 2018). "Wage dispersion and search behavior: The importance of nonwage job values". In: *Journal of Political Economy* 126.4, pp. 1594–1637.
- Hamermesh, Daniel S (1999). "The timing of work over time". In: *The Economic Journal* 109.452, pp. 37–66.

- Hausman, Jerry A. (Nov. 1978). "Specification Tests in Econometrics". In: *Econometrica* 46.6, pp. 1251–1271.
- Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska (2022). "National wage setting". In: University of Chicago, Becker Friedman Institute for Economics Working Paper 2022-150.
- Hintermann, Beat, Anna Alberini, and Anil Markandya (2010). "Estimating the value of safety with labour market data: are the results trustworthy?" In: *Applied Economics* 42.9, pp. 1085–1100.
- Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed (Oct. 1998). "Hedonic Wages and Labor Market Search". In: *Journal of Labor Economics* 16.4, pp. 815– 847.
- Johnson, Matthew S, David I Levine, and Michael W Toffel (Oct. 2023). "Improving regulatory effectiveness through better targeting: Evidence from OSHA". In: *American Economic Journal: Applied Economics* 15.4, pp. 30–67.
- Johnston, Andrew C (2020). "Preferences, selection, and the structure of teacher pay". In: *Available at SSRN 3532779*.
- Kline, Patrick (Dec. 2024). "Firm Wage Effects". In: ed. by Christian Dustmann and Thomas Lemieux. Vol. 5. Handbook of Labor Economics. Amsterdam, NL: Elsevier. Chap. 2, pp. 115–182.
- Kniesner, Thomas J and John D Leeth (1995). "Abolishing OSHA". In: *Regulation* 18, pp. 46–56.
- Kniesner, Thomas J, W Kip Viscusi, Christopher Woock, and James P Ziliak (Feb. 2012). "The value of a statistical life: Evidence from panel data". In: *Review of Economics and Statistics* 94.1, pp. 74–87.
- Kostiuck, Peter F. (Oct. 1990). "Compensating Differentials for Shift Work". In: *Journal of Political Economy* 98.5 Part 1, pp. 1054–1075.
- Krueger, Alan B and David Schkade (Fall 2008). "Sorting in the labor market: do gregarious workers flock to interactive jobs?" In: *Journal of Human Resources* 43.4, pp. 859–883.
- Krueger, Alan B and Lawrence H Summers (1988). "Efficiency wages and the interindustry wage structure". In: *Econometrica: Journal of the Econometric Society*, pp. 259–293.
- Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen A Woodbury (Apr. 2023a). "Do Firm Effects Drift? Evidence from Washington Administrative Data". In: *Journal of Econometrics* 233.2, pp. 375–395.
- (May 2023b). "Work hours mismatch". In: NBER Working Paper Series 31205.

- Lagos, Lorenzo (May 2024). "Union Bargaining Power and the Amenity-Wage Tradeoff". In: *IZA Discussion Paper Series* 17034.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (Jan. 2022). "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market". In: American Economic Review 112.1, pp. 169–212.
- Lang, Kevin and Sumon Majumdar (Nov. 2004). "The pricing of job characteristics when markets do not clear: theory and policy implications". In: *International Economic Review* 45.4, pp. 1111–1128.
- Lavetti, Kurt (Summer 2023). "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications". In: *Journal of Economic Perspectives* 37.3, pp. 189–212.
- Lavetti, Kurt and Ian M. Schmutte (2018). *Estimating Compensating Wage Differentials* with Endogenous Job Mobility. Working Paper.
- Lee, Jonathan M and Laura O Taylor (Nov. 2019). "Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life". In: *American Economic Journal: Economic Policy* 11.4, pp. 350–374.
- Levine, David I., Michael W. Toeffel, and Matthew S. Johnson. (May 2012). "Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss". In: *Science* 336.6083, pp. 907–911.
- Levitt, Larry and Gary Claxton (Oct. 2011). *What the Actuarial Values in the Affordable Care Act Mean.* Tech. rep. 8177. Menlo Park, CA: Kaiser Family Foundation.
- Lewandowski, Piotr, Katarzyna Lipowska, and Mateusz Smoter (2022). "Working from Home during a Pandemic–A Discrete Choice Experiment in Poland". In.
- Lewis, H. Gregg (Aug. 1969). "Interes del Empleador en las Horas de Trabajo del Empleado". In: *Cuadernos de Economía* 6.18, pp. 38–54.
- Luce, R Duncan (1959). *Individual choice behavior*. New York, NY: John Wiley & Sons.
- Maestas, Nicole, Kathleen J. Mullen, David Powell, Till von Wachter, and Jeffrey B. Wenger (July 2023). "The Value of Working Conditions in the United States and the Implications for the Structure of Wages". In: *American Economic Review* 113.7, pp. 2007–2047.
- Mandeville, Kate L, Mylene Lagarde, and Kara Hanson (2014). "The use of discrete choice experiments to inform health workforce policy: a systematic review". In: *BMC health services research* 14, pp. 1–14.
- Manning, Alan (Apr. 2003). "The real thin theory: monopsony in modern labour markets". In: *Labour Economics* 10.2, pp. 105–131.

- Mas, Alexandre (Jan. 2008). "Labour unrest and the quality of production: Evidence from the construction equipment resale market". In: *Review of Economic Studies* 75.1, pp. 229–258.
- Mas, Alexandre and Amanda Pallais (Dec. 2017). "Valuing Alternative Work Arrangements". In: *American Economic Review* 107.12, pp. 3722–3759.
- (2019). "Labor supply and the value of non-work time: Experimental estimates from the field". In: *American Economic Review: Insights* 1.1, pp. 111–126.
- (Aug. 2020). "Alternative Work Arrangements". In: Annual Review of Economics 12, pp. 631–658.
- Monaco, Kristen and Brooks Pierce (July 2015). "Compensation inequality: evidence from the National Compensation Survey". In: *Monthly Labor Review* 138.7.
- Morchio, Iacopo and Christian Moser (May 2024). "The Gender Pay Gap: Micro Sources and Macro Consequences". In: *NBER Working Paper Series*.
- Mortensen, Dale T. (1986). "Job Search and Labor Market Analysis". In: ed. by Orley C. Ashenfelter and Richard Layard. Vol. 2. Handbook of Labor Economics. Amsterdam, NL: Elsevier. Chap. 15, pp. 849–919.
- Nagler, Markus, Johannes Rincke, and Erwin Winkler (Sept. 2024). "Working from home, commuting, and gender". In: *Journal of Population Economics* 37.3.
- Negahban, Sahand, Sewoong Oh, and Devavrat Shah (Jan. 2017). "Rank Centrality: Ranking from Pairwise Comparisons". In: *Operations Research* 65.1, pp. 266–287.
- Nekoei, Arash (Jan. 2022). "Will Markets Provide Humane Jobs? A Hypothesis". In: *CESifo Working Paper Series*.
- Non, Arjan, Ingrid Rohde, Andries de Grip, and Thomas Dohmen (2022). "Mission of the company, prosocial attitudes and job preferences: A discrete choice experiment". In: *Labour Economics* 74, p. 102087.
- Office of Management and Budget (Sept. 2003). *Regulatory Analysis*. OMB Circular No. A-4.
- Ouimet, Paige and Geoffrey Tate (July 2023). "Firms with Benefits? Nonwage Compensation and Implications for Firms and Labor Markets". In: *NBER Working Paper Series*.
- Oyer, Paul (2008). "Salary or benefits?" In: Work, earnings and other aspects of the employment relation. Ed. by Solomon W. Polachek and Konstantinos Tatsiramos. Vol. 28. Emerald Group Publishing Limited. Chap. 13, pp. 429–467.
- Page, Lawrence (1999). *The PageRank citation ranking: Bringing order to the web*. Tech. rep. Technical Report.

- Park, Sangyoon (July 2016). "Socializing at Work: Evidence from a Field Experiment with Manufacturing Workers". In: *American Economic Journal: Applied Economics* 11.3, pp. 424–455.
- Pierce, Brooks (Nov. 2001). "Compensation Inequality". In: *Quarterly Journal of Economics* 116.4, p. 2001.
- Quach, Simon (2025). "Wage Hysteresis and Entitlement Effects: The Persistent Impacts of a Temporary Overtime Policy". In: *The Quarterly Journal of Economics*, qjaf008.
- Roback, Jennifer (Dec. 1982). "Wages, Rents, and the Quality of Life". In: *Journal of Political Economy* 90.6, pp. 1257–1278.
- Rosen, Sherwin (Jan. 1974). "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition". In: *Journal of Political Economy*, 82.1, pp. 34–55.
- (1986). "The Theory of Equalizing Differences". In: ed. by Orley C. Ashenfelter and Richard Layard. Vol. 1. Handbook of Labor Economics. Amsterdam, NL: Elsevier. Chap. 12, pp. 641–692.
- Roussille, Nina and Benjamin Scuderi (July 2023). "Bidding for Talent: A Test of Conduct in a High-Wage Labor Market". In: *IZA Discussion Paper Series* 16352.
- Salop, Steven C (1979). "A model of the natural rate of unemployment". In: *The American Economic Review* 69.1, pp. 117–125.
- Seo, Sukyong and Joanne Spetz (Aug. 2013). "Demand for temporary agency nurses and nursing shortages". In: *INQUIRY: The Journal of Health Care Organization*, *Provision, and Financing* 50.3, pp. 216–228.
- Shapiro, Carl and Joseph E Stiglitz (1984). "Equilibrium unemployment as a worker discipline device". In: *The American economic review* 74.3, pp. 433–444.
- Slichter, Sumner H. (Feb. 1950). "Notes on the Structure of Wages". In: *The Review of Economics and Statistics* 32.1, pp. 80–91.
- Smith, Adam (1776 [2008]). An Inquiry into the Nature and Causes of the Wealth of Nations. Ed. by Kathryn Sutherland. Oxford, UK: Oxford University Press. ISBN: 9780199535927.
- Sockin, Jason (July 2022). "Show Me the Amenity: Are Higher-Paying Firms Better All Around?" In: *CESifo Working Paper Series*.
- Sorkin, Isaac (Aug. 2018). "Ranking Firms Using Revealed Preference". In: *Quarterly Journal of Economics* 133.3, pp. 1331–1393.
- Stern, Scott (June 2004). "Do Scientists Pay to be Scientists?" In: *Management Science* 50.6, pp. 835–853.
- Sullivan, Paul and Ted To (2014). "Search and nonwage job characteristics". In: *Journal of Human Resources* 49.2, pp. 472–507.

- Summers, Lawrence H. (May 1989). "Some Simple Economics of Mandated Benefits". In: *American Economic Review* 79.2, pp. 177–183.
- Syverson, Chad (June 2011). "What Determines Productivity?" In: *Journal of Economic Literature* 49.2, pp. 326–365.
- Taber, Christopher and Rune Vejlin (2020). "Estimation of a roy/search/compensating differential model of the labor market". In: *Econometrica* 88.3, pp. 1031–1069.
- Tsao, Carolyn C. (Jan. 2025). "It's Not (Just) About the Money: Pay and the Value of Working Conditions in Teaching". PhD thesis. Princeton, N.J.: Princeton University.
- U.S. Department of Transportation (Aug. 2016). *Revised Departmental Guidance 2016: Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses.* Memorandum.
- Van Ommeren, Jos, Gerard J. Van Den Berg, and Cees Gorter (2000). "Estimating the Marginal Willingness to Pay for Commuting". In: *Journal of Regional Science* 40.3, pp. 541–563.
- Villanueva, Ernesto (July 2007). "Estimating Compensating Wage Differentials Using Voluntary Job Changes: Evidence from Germany". In: *Industrial and Labor Relations Review* 60.4, pp. 544–561.
- Viscusi, W Kip (Winter 2015). "The role of publication selection bias in estimates of the value of a statistical life". In: *American Journal of Health Economics* 1.1, pp. 27–52.
- Wissmann, Daniel (Jan. 2022). "Finally a Smoking Gun? Compensating Differentials and the Introduction of Smoking Bans". In: *American Economic Journal: Applied Economics* 14.1, pp. 75–106.
- Wiswall, Matthew and Basit Zafar (Feb. 2018). "Preference for the Workplace, Investment in Human Capital, and Gender". In: *The Quarterly Journal of Economics* 133.1, pp. 457–507.
- Woodbury, Stephen A. (Mar. 1983). "Substitution between Wage and Nonwage Benefits". In: *The American Economic Review* 73.1, pp. 166–182.
- Woodbury, Stephen A. and Daniel S. Hamermesh (May 1992). "Taxes, Fringe Benefits and Faculty". In: *Review of Economics and Statistics* 74.2, pp. 287–296.
- Woodbury, Stephen A. and Wei-Jang Huang (1991). *The Tax Treatment of Fringe Benefits*. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.