Journal of Economic Perspectives—Volume 37, Number 3—Summer 2023—Pages 189–212

Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications

Kurt Lavetti

How do workers choose which jobs to accept? Although earnings are one important factor in this decision, workers typically also consider many other job characteristics. They may be willing to accept jobs that pay less but offer flexible hours, health insurance, or shorter commutes. Dangerous or dirty jobs like coal mining or trash collection may have to offer higher pay to entice workers to accept the undesirable characteristics of the job. Of course, people may disagree about whether a particular job characteristic is good or bad. Some people prefer working at a desk, while others prefer being outside or doing physically active work.

Conceptually, choosing a job can be thought of as a worker selling their services in the labor market and simultaneously buying amenities (which can be positive or negative) from their employer. The observed wage rate combines the value of a worker's time and the implicit prices of all amenities. Quantifying the tradeoffs between earnings and job characteristics, also called "compensating wage differentials," is of fundamental importance for understanding labor market equilibria and wage dispersion.

To illustrate, consider the differences in average wages and education levels across occupations shown in Table 1. Notice that although bakers have six times higher college graduation rates than butchers, they earn about 8 percent less on average. This may be, in part, because most people would prefer, all else equal, to work in a redolent bakery full of fresh bread than in a butchery, so butchers must be paid extra for the undesirable characteristics of their job. And although accountants

■ *Kurt Lavetti is Associate Professor of Economics, The Ohio State University, Columbus, Ohio and a Research Associate at the National Bureau of Economic Research. His email address is lavetti.1@osu.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at https://doi.org/10.1257/jep.37.3.189.

Table 1
Average Education and Annual Full-Time Earnings of Selected Occupations

Occupation	*High school degree*	*College degree*	*Annual earnings*
Butcher	0.70	0.03	$36,150
Baker	0.84	0.18	$33,570
Accountant	1.00	0.79	$72,020
Explosives handler	0.92	0.00	$100,120
Forest conservationist	0.97	0.39	$36,760
Agricultural inspector	0.94	0.39	$49,460

Source: IPUMS March CPS data (Flood et al. 2022).
Note: Averages are calculated using March CPS data from 2012–2022. Columns 1 and 2 report the average share of workers by occupation with at least a high school degree and at least a college degree, respectively. Column 3 reports average annual earnings of a full-time equivalent worker by occupation using 2020 dollars.

are highly educated, they earn 28 percent less than explosives handlers, who have low college completion rates but are exposed to physical danger. Finally, forest conservationists earn 26 percent less than agricultural inspectors despite similar education levels. Workers who choose forest conservation jobs may be willing to accept lower salaries in exchange for the opportunity to work outdoors, frequently in scenic parks, and also to feel good about contributing to environmental stewardship. These are just examples, but any feature of a job that workers have preferences over could potentially have an associated compensating wage differential.

Although the concept of a compensating differential is straightforward and intuitive, its simplicity belies how difficult it is to empirically quantify compensating differentials in real-world labor markets. Although Adam Smith (1776) wrote about the tradeoffs between earnings and job characteristics,[1] it would take nearly 200 years until Sherwin Rosen (1974; 1986) formalized a theoretical and empirical framework for studying compensating differentials. In this article, I begin with an overview of the Rosen (1974; 1986) model, which emphasizes that compensating wage differentials result from workers with different preferences for amenities sorting between firms with different costs of providing amenities.

I then present a chronology of the empirical approaches used to estimate compensating differentials, highlighting the new lessons learned as data quality and methods advanced over time and how these advances in turn revealed new challenges and setbacks. I begin with basic cross-sectional wage models, and then discuss how panel data models were used to alleviate bias caused by unobserved worker skills. Attempts to extend the theory of compensating wage differentials to markets with imperfect competition or search frictions led to grave concerns about

[1] In *The Wealth of Nations* (Book X, Part I), Adam Smith (1776) writes "[T]he wages of labour vary with the ease or hardship, the cleanliness or dirtiness, the honourableness or dishonourableness of the employment. A . . . blacksmith . . . seldom earns so much . . . as a [coal miner] . . . . His work is not quite so dirty, is less dangerous, and is carried on in daylight, and above ground. . . ."

the applicability of the Rosen model to realistic labor market settings and sowed doubt about the reliability of empirical estimates. Recent progress in responding to these concerns has involved using newly available matched worker-firm data and settings with quasi-random variation in levels of job amenities.

Empirical estimates of compensating differentials are important for informing and designing a broad range of public policies. One example is that estimates of the compensating differential for the risk of death on a job are used to calculate the value of statistical life, which in turn is widely used in cost-benefit analyses by government regulators. Compensating differentials are also important for interpreting and measuring earnings disparities, which inform public policies aimed at reducing inequality. For example, Sorkin (2018) estimates that compensating wage differentials explain at least 15 percent of earnings inequality in the United States, while Taber and Vejlin (2020) estimate that they explain up to 26 percent of wage inequality in Denmark. To put the importance of compensating differentials in a more direct perspective, the authors estimate that if all amenities were eliminated from all jobs, total wages would increase by about 18 percent. In other words, compensating differentials are roughly as important for aggregate economic activity as the entire healthcare sector is for the US economy.

The Rosen Model: Parameters and Interpretations

The Rosen (1986) model provides a theoretical framework for understanding tradeoffs between wages and amenities in labor markets. Consider a competitive labor market in which workers choose between jobs that offer different levels of wages, W, and amenities, A, to maximize utility $u_i(W, A)$. A includes both good and bad job characteristics, where a reduction in disamenities can be expressed as an increase in amenities, and vice versa. The subscript on u_i indicates that different workers, denoted by i, have different preferences for amenities relative to wages.

In this model, all workers are assumed to have the same productivity—the model abstracts from factors other than A that might cause wages to vary, making it important to control these other factors in empirical estimation. Workers choose whether to sell their services in the labor market at wage rate W and also choose how much A to buy, where the price of A is an implicit reduction in the wage rate. In this sense, a worker's decision to accept a job can be viewed as a simultaneous selling and buying choice. The observed wage is the sum of these two transactions that are tied together but conceptually distinct.

Let's begin by focusing on the worker's side of the labor market in a basic discrete version of the model, in which jobs offer either zero amenities or fixed levels A_1, A_2, or A_3. Examples of discrete amenities include the number of days per week employees can work from home or weeks of vacation time per year. Workers who choose to accept a job with zero amenities earn the market wage, denoted W_0. Jobs with more amenities can offer lower wages and still attract workers, creating a downward-sloping relationship between equilibrium wages and amenities, as shown in the

Figure 1.
Willingness to Pay for Amenities and Job Choices

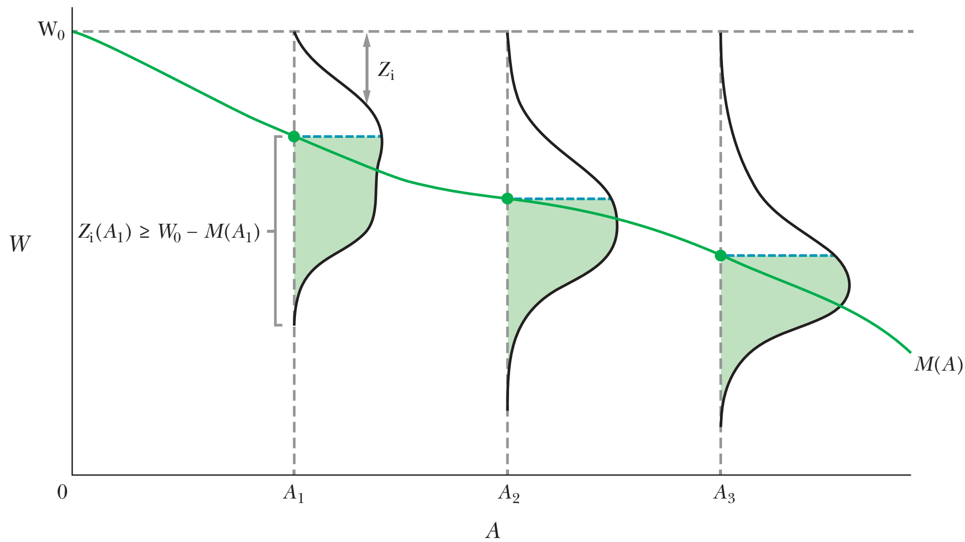

Source: Figure created by author.
Note: The horizontal axis shows levels of amenities A. The vertical axis shows wages W. W_0 is the wage for a job with zero amenities. For each discrete level of amenities A_1, A_2, and A_3, the Z_i curve shows the distribution of preferences for the given amenity level across individuals. The $M(A)$ function shows the downward-sloping relationship between wages and amenities in equilibrium. All the workers in the shaded region, for whom $Z_i(A) > W_0 - M(A)$, value amenities by more than the market reduction in wages required to obtain them.

$M(A)$ function in Figure 1. The vertical distance between W_0 and the $M(A)$ function is how much of their potential wage workers give up in exchange for amenities.

Although there is a single market wage for jobs that provide A_1, workers have differing preferences. Imagine that each worker was asked: what is the maximum reduction in wages, Z_i, you would accept to obtain a job with amenity level A_1? Because workers have heterogeneous preferences, there is a distribution of Z_is. The solid black line in Figure 1 shows this distribution (oriented along the vertical axis). Workers who value amenities the most have high Z_is, indicating a high willingness to give up wages in exchange for amenities. These workers are located on the bottom tail of the distribution, furthest from W_0. Those who do not value amenities at all have $Z_i = 0$ and are located at the top end of the black line, where $W = W_0$. As the figure suggests, the distributions of Z_is could be arbitrary, but for simplicity consider the case where $Z_i \geq 0$, as the equilibrium is more complicated if workers disagree about whether A is good or bad. The $M(A)$ function splits workers into two groups: those with small Z_is are unwilling to accept the market wage $M(A)$ and prefer to take the job with zero amenities, while those with large Z_is (the shaded region) prefer to take a job with amenities.

A key feature of this model is that workers who hold jobs with amenities earn (weakly) positive surplus, or rents, from these jobs. All the workers in the shaded region, for whom $Z_i(A) > W_0 - M(A)$, value amenities by more than the market reduction in wages required to obtain them. Because of rents, if two workers who chose jobs with different levels of A were forced to switch jobs with each other, total surplus would generally fall. This is unlike equilibria in most economic models of competitive markets, where the matching between particular buyers and sellers is not relevant to total welfare. The welfare effects of rents in the Rosen model are similar to the effects of bundling jobs with purchases of goods that yield positive consumer surplus.

The presence of rents also has implications that extend beyond the Rosen model. For example, in models of equilibrium wages that are based on negotiations between workers and firms, rents can impact the outcomes of negotiations, affecting wages and job mobility patterns.

So far, the $M(A)$ function has been taken as given, but what determines this function? Answering this question requires considering labor demand and the employer side of the problem. The original Rosen (1986) model considers the case of firms with linear technology, where output x is produced for sale, labor L is a production input, and $\alpha_j(A)$ is a linear productivity parameter. The productivity parameter may differ across firms, denoted by j. In order to provide amenities, firms must give up productivity; that is, the cost of providing A can be thought of as foregone units of output x. As an example, consider workplace safety on a manufacturing assembly line: increasing safety may require slowing the pace of the assembly line, which decreases productivity. To choose the optimal level of amenities A, the firm equates the marginal benefit and marginal cost of providing A. The marginal benefit comes from firms paying lower wages when they offer A, and the marginal cost comes from foregone productivity. Firms differ in the technology for providing amenities; those with lower marginal costs choose to offer higher levels of A, all else equal.[2]

Firms face tradeoffs between offering higher wages or more amenities, and have isoprofit curves that describe the (W, A) pairs that hold profits fixed. Figure 2 depicts these isoprofit curves, denoted $\phi_1 - \phi_4$, for four firms that face different marginal costs of providing A. Firm ϕ_1 faces a high cost, so its isoprofit functions begin to decline steeply at low levels of A. Firms $\phi_2 - \phi_4$ have different technology functions that allow them to provide amenities at a relatively lower cost.

Workers' preferences can also be expressed by indifference curves over wages and amenities. For each worker, the optimal wage-amenity pair is the point along the $M(A)$ function that maximizes utility. As Figure 2 depicts, workers with preferences described by indifference curve θ_1 maximize utility by choosing the point on $M(A)$ that is tangent to their indifference curve. Workers with different preferences

[2] Although this assumption that amenity levels affect firm productivity is used in the original Rosen (1986) model, it is straightforward to change this assumption and consider the case in which A is costly but has no direct effect on productivity.

Figure 2
Sorting and Matching in the Labor Market

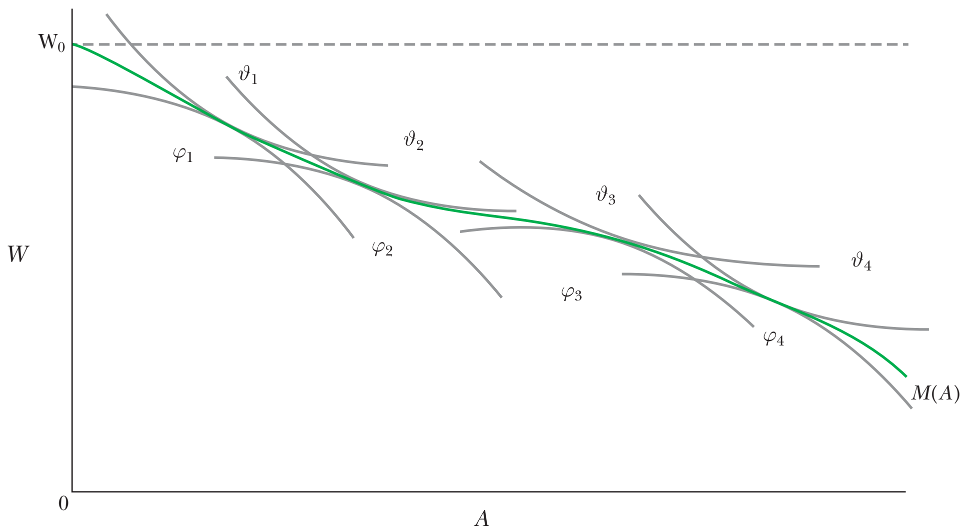

Source: Figure created by author.
Note: The horizontal axis shows amenity levels, A. The vertical axis shows wage levels, W. Four isoprofit curves, denoted $\phi_1 - \phi_4$, illustrate four firms that face different marginal costs of providing A. Four indifference curves, denoted $\theta_1 - \theta_4$, illustrate the preferences of four workers over tradeoffs between wages and amenities. The market equalizing differences curve, $M(A)$, is traced out by the set of tangency points between workers' indifference curves and firms' isoprofit functions.

sort along the $M(A)$ function according to their relative preferences for wages and amenities.

Putting both sides of the market together, the market equalizing differences curve, $M(A)$, is defined by the set of tangency points between workers' indifference curves and firms' isoprofit functions. In equilibrium, workers with different preferences for amenities are matched to firms with different marginal costs of providing amenities. The sorting and matching process that generates $M(A)$ is an essential component of the equilibrium. The shape of the $M(A)$ function depends on many factors, including the distribution of preferences, the distribution of firm technology parameters, output prices, the relative supply of workers at each level of A, and the relative demand for workers at each level of A. $M(A)$ is not generally a linear function.

In this literature, researchers have defined the idea of "compensating wage differentials" in various ways. I follow Rosen (1986) in defining the terms "compensating wage differential" and "equalizing wage differential" synonymously as the slope of the $M(A)$ market equalizing differences curve—that is, the rate at which market wages change as amenity levels change.

In this equilibrium, three classes of parameters are frequently objects of interest for empirical researchers and policymakers: (1) parameters that characterize workers' preferences and their marginal willingness to "pay" for amenities via lower wages; (2) parameters that characterize firms' costs of providing an amenity; and (3) parameters that characterize the $M(A)$ function or its slope, the market compensating wage differential. The empirical approaches to estimating these classes of parameters differ, and conflating these three conceptual objects has been a source of confusion for empirical interpretation and policy applications.

In particular, the $M(A)$ function is generally neither a representation of workers' preferences nor firms' technology—it arises from the sorting process that matches workers to firms. The market compensating wage differential is a local measure of preferences for the *marginal worker* whose indifference curve is tangent to the $M(A)$ function at a certain level of amenities. Among all workers who accept a job with a given amenity level, the marginal worker is the worker with the lowest willingness to pay. Therefore, the slope of the $M(A)$ curve provides little information about the preferences of an average worker. However, it does provide a lower bound on the preferences of inframarginal workers.

As an example, consider the case of jobs that require traveling. Some workers may prefer the opportunity to travel for work, while others may view work-related travel as disruptive to their lives.[3] The marginal worker might be someone with preferences somewhere in the middle. Because the $M(A)$ function is based on the preferences of only the marginal worker, the market compensating wage differential for jobs that require travel could potentially be zero even if most workers have strong preferences in one direction or the other.

In addition, the $M(A)$ function can shift for reasons unrelated to either the preferences of workers or the technology of firms. A shock to the output price of a good can move firms' isoprofit functions, shifting the profit-maximizing tradeoff between wages and amenities. For example, Charles et al. (2022) find that a 10 percent shock to output prices leads to a 1.5 percent increase in injury rates on average. Even if all workers' preferences remain the same, labor demand shifts may change how workers are sorted into jobs and thus shift the $M(A)$ function. With a single cross-sectional dataset, it is not typically possible to identify the full distribution of workers' preferences, or even *average* preferences, without imposing further assumptions, such as an assumption on the preference distribution (Ashenfelter 2006; Rosen 1986).[4]

Some special cases of the Rosen model can simplify estimation (Hwang, Mortensen, and Reed 1998; Rosen 1974). If firms have the same technology functions, workers with heterogeneous preferences sort along a single common isoprofit function. In this case, the market equalizing difference function is equivalent to the isoprofit function. Variation in labor supply can be used to estimate firm technology parameters. Similarly, if workers share common preferences, then changes in firm

[3] For related work on how workers sort differently across jobs after having children, see Hotz, Johansson, and Karimi (2018), who study shifts towards jobs with more family friendly amenities such hours flexibility.
[4] There are of course other approaches to estimating workers' preferences that do not rely on equilibrium labor market data. See Mas and Pallais (2017) for an example.

technology can be used to trace out the slope of indifference curves, identifying preferences.

A related scenario, an important special case for empirical researchers, occurs when firms share a common isoprofit function and purchase amenities in a competitive intermediate market instead of producing them. For example, firms may purchase health insurance for workers or contribute to retirement savings accounts. This scenario differs from the model depicted in the above graphs if providing amenities does not impact productivity and firms all face the same market price of amenities. Again, if firms share a common isoprofit function in a perfectly competitive labor market, workers with different preferences sort along this function and the isoprofit function is equivalent to $M(A)$. The $M(A)$ function may be linear if all firms face the same constant unit cost per worker for the amenity in a competitive labor market (and there are no tax-related distortions). In this case, the market-compensating wage differential is simply the market cost of the amenity, which may be observable to researchers. Of course, this theoretical result relies upon strong assumptions, including perfectly competitive labor and amenity markets, which may not hold in real-world empirical settings like markets for health insurance.

A Chronology of Empirical Methods, Estimation Challenges, and Limitations

I begin by presenting a cross-sectional approach taken in many earlier studies and then discuss adjustments that researchers made as panel data became more widely available. These panel-based approaches are designed to account for unobserved differences across workers and employers and imperfectly competitive labor markets. Finally, I discuss recent methods combining panel data with natural experimental designs involving exogenous amenities variation.

Cross-Sectional Wage Models

An extensive literature has sought to estimate compensating wage differentials by using a framework built on an ordinary least squares regression like:

$$W_i = X_i\beta + A_i\gamma + \varepsilon_i,$$

where W_i is the log wage of worker i; X_i is a vector of observed worker, job, and employer wage shifters, in addition to, ideally, any other market factors that may shift wages for reasons unrelated to amenities, such as time and location; and A_i is a vector of observed job amenities.[5] The aspiration behind this approach is to include in X_i all of the factors that systematically shift wages. For example, X_i should include worker skills or ability measures that affect worker productivity so that conditional

[5] For readers who wish to learn more about this cross-sectional literature, some helpful starting points include Viscusi and Aldy (2003), Duncan (1976), and Brown (1980).

on X_i workers are productively homogeneous. Under the assumption that X_i contains a sufficiently rich set of control variables, γ provides information about the average compensating wage differential for each amenity A. In practice, there are many challenges to implementing this approach. Here, I focus on a few of the most important and pervasive estimation and interpretation challenges, although a complete account is beyond the scope of this article.

In an example of this approach, Garen (1988) studies the relationship between log wages and the risk of injury on the job and estimates the amenity coefficient γ to be 0.0024. This result suggests that if a job increases a worker's annual risk of death by 1 in 100,000, then wages are on average higher by about 0.24 percent. For reference, a typical US manufacturing job has a fatal injury rate of about 3.5 deaths per 100,000 worker-years.[6] Therefore, the author's estimate of γ suggests that average manufacturing wages are about 0.84 percent higher than they would be if all fatal injury risks could be (hypothetically) eliminated.

The basic regression version of the cross-sectional model assumes that the market-equalizing function is linear (with slope γ), thus ruling out the possibility of heterogeneity in compensating wage differentials. However, it is generally improbable in the Rosen model that all the tangency points between indifference curves and isoprofit functions will happen to lie on a straight line (Ekeland, Heckman, and Nesheim 2004). Thus, it is typically good practice to evaluate the linearity assumption relative to a more flexible function that may be nonlinear in A, although researchers frequently have limited statistical power to reject one functional form in favor of another.

Another challenge with the cross-sectional approach is that it is impractical to measure all of the dimensions of worker skills or abilities that influence wages and include them in X_i. Any unobserved components of ability are likely to cause endogeneity bias when estimating γ because high-earning workers tend to prefer jobs with more amenities, creating a positive correlation between ε and A.

In a classic example, Lucas (1977) estimates a cross-sectional wage model that controls for age, gender, schooling, and union membership and estimates compensating differentials for a set of amenities, including workplace hazards, whether the job requires repetitive tasks, is physically demanding, or requires supervising workers (among other amenities). Lucas finds that among workers with a high school degree, jobs that require highly repetitive tasks pay around 10–25 percent higher hourly wages. An instructive if perplexing finding, however, is that the coefficient on physically demanding jobs is negative—that is, it suggests that workers are willing to sacrifice wages to obtain physically demanding jobs. As Lucas (p. 218) explains: "[I]t is highly probable this effect may be explained by the omission of some ability, other than schooling, which is possessed by people in sedentary occupations." Although Lucas does not address this problem of unobserved ability and

[6] This estimate is based on the Census of Fatal Occupation Injuries, available at https://www.bls.gov/iif/fatal-injuries-tables.htm.

Figure 3
Unobserved Skills as a Source of Bias

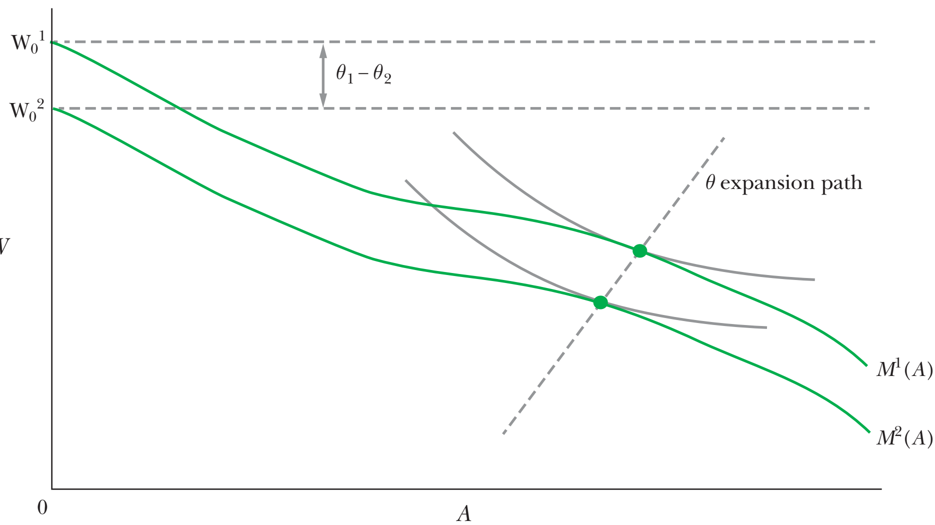

Source: Figure created by author.
Note: The horizontal axis shows amenity levels, A. The vertical axis shows wage levels, W. There are two types of workers: high productivity (θ_1) or low productivity (θ_2). High-productivity workers can access jobs along the equalizing differences curve $M^1(A)$, while low-productivity workers can access jobs along the lower curve, $M^2(A)$. If the amenity A is a normal good, then the utility-maximizing choice of wages and amenities (W, A) will lie along an upward-sloping path as worker productivity rises.

omitted variables, his speculation that unobserved ability may explain counterintuitive findings in cross-sectional studies was echoed in subsequent literature.

To gain some intuition about such counterintuitive results, suppose the assumption that workers are productively homogeneous does not hold. Instead, there are two types of workers with either high productivity (θ_1) or low productivity (θ_2). Productivity is observed by firms but unobserved by the econometrician. As Figure 3 depicts, high-productivity workers can access jobs along the equalizing differences curve $M^1(A)$, while low-productivity workers can access jobs along the lower curve, $M^2(A)$. If the amenity A is a normal good, income effects will lead workers with high productivity to use some of their additional earnings potential to buy more A. Thus, as productivity increases (holding preferences fixed), the utility-maximizing choice of wages and amenities (W, A) will lie along an upward-sloping path as worker productivity rises. The same intuition extends to workers with different preferences—at each point along $M(A)$ there is an upward-sloping θ expansion path.

A researcher using data from this labor market faces the problem that there are two types of variation in the data—variation due to preference-based sorting along the negatively-sloped market equalizing differences curve and variation due to ability differences along the positively-sloped expansion path. In addition, worker productivity θ_i is correlated with amenities A_i under the assumption that amenities are normal goods. For example, if one runs a regression of log wages on job amenities only, without controlling for any differences in education or human capital, wages are usually positively correlated with amenities, because highly-skilled workers disproportionately hold good, high-paying jobs that offer many amenities.[7] Therefore, failing to control for productivity shifters in θ_i will introduce bias in the γ coefficient. If A is a desirable amenity, the bias is usually positive, and if A is an undesirable disamenity, the bias is usually negative.[8]

Unobserved productivity or ability is just one example of an omitted variable that may introduce bias in the cross-sectional model. It is important to note that the same intuition behind this source of bias extends broadly to many different forms of systematic variation in earnings besides worker productivity. Researchers must also consider many other sources of wage dispersion among workers that might be correlated with the amenities A_i and/or the controls in X_i. For example, in wage models with search frictions, more experienced workers often earn systematically higher wages. Because these workers have higher earnings potential, they may choose to purchase more amenities A, creating a correlation between experience, the degree of friction in the labor market, and amenities. For this reason, having a thorough model of wages that incorporates many sources of wage dispersion can help alleviate bias in the estimation of compensating wage differentials. In practice, however, no matter how detailed the available data, it is difficult in a cross-sectional model to completely account for all of the factors that may influence wages, including ability, unobserved amenities, labor market frictions, efficiency wages, rent-sharing, or discrimination, among others—all of which are potential sources of estimation bias.

In some cases, researchers have reason to believe that compensating differentials vary with observed characteristics of workers or jobs; for example, workers with young children may be highly averse to the risk of a severe injury. Researchers may wish to quantify the differences in compensating differentials across such groups. However, estimating heterogeneity in compensating differentials requires a strong word of caution, as Thaler and Rosen (1976) explained. In their study, the authors use cross-sectional data to estimate the relationship between wages and the risk of injury on the job using an ordinary least squares regression of the form:

$$W_i = X_i\beta + A_i\gamma_1 + A_i s_i \gamma_2 + \varepsilon_i,$$

where A_i is a measure of injury risk, s_i are worker characteristics that interact with injury risk to affect wages, and the remaining variables are defined above. An example of s_i that the authors consider is age. If a severe injury to a younger worker

[7] In some scenarios, the selection problem is more complex. See DeLeire, Khan, and Timmins (2013) for a model of compensating wage differentials with Roy-style sorting in which the returns to human capital differ by occupation.

[8] However, the direction of bias is not straightforward if the omitted productivity component θ_i is multidimensional, or if A is multidimensional.

causes a greater decline in lifetime future consumption, one may hypothesize that compensating wage differentials are larger for younger workers. Another possibility is that individual risk preferences may change with age for reasons other than lifetime future consumption.

This approach may appear to be a straightforward extension of the basic ordinary least squares cross-sectional specification; however, as Thaler and Rosen (1976) cautioned, the interpretation of estimates from specifications of this form is actually a "rather thorny issue." The reason is that in the Rosen model, there can only be one single market equalizing differences function, $M(A)$, for productively equivalent workers. However, an empirical specification that also includes the heterogeneity term $A_i s_i \gamma_2$ will shift the market equalizing differences function vertically for workers with different values of s_i. Therefore, an empirical model that includes shifters in the market equalizing differences function to capture preference heterogeneity is fundamentally inconsistent with the Rosen model.

However, including shifters for combinations of personal traits and amenities like $A_i s_i \gamma_2$ in an empirical model is not *always* problematic. If the researcher believes that workers with different values of s_i have different productivity, *and* the differences in productivity scale with the level of amenities, then including this term in the model is an appropriate—and, in general, necessary—way to account for productive heterogeneity. Just as the omission of any productivity shifter (θ_i) from the model can generally cause endogeneity bias, so too can the omission of factors that shift productivity in ways that depend on amenity levels. In Thaler and Rosen's (1976) model s_i captures observed factors that shift whether workers are able to be productive in the presence of risks—say, whether workers have "nerves of steel." If s_i includes age, the model imposes a specification assumption that older workers with more experience have a relative advantage in performing tasks in the presence of physical hazards, so that productivity differences by age scale with risk A_i. When estimating this type of specification, researchers should use caution to clearly separate how functional form assumptions related to the equalizing differences function differ from assumptions related to how the worker productivity component might interact with amenities.[9]

This point also raises the general question: how might researchers describe heterogeneity in compensating wage differentials? Instead of using a shifter term like $A_i s_i \gamma_2$, it is first necessary to relax the assumption that log wages change linearly with A_i and instead allow γ to vary with A_i. After estimating a model of this form, one can quantify group-level differences in compensating wage differentials by estimating the average marginal effects of changes in amenities on wages for different groups of workers. If workers with different characteristics sort to different locations on this nonlinear equalizing differences function, the average slope of the

[9] Hersch (2011) also considers a model similar to this in which the wage effects of a disamenity, workplace sexual harassment rates, operate through both a productivity channel and a compensating wage differential channel.

equalizing differences function may differ across groups, providing information about heterogeneity in compensating wage differentials.

Panel Models: Unobserved Skills and the Ability Bias Puzzle

To address unobserved differences in worker skills, researchers can use panel data that track the same workers over time and in this way control for some traits unobservable in standard data, such as ability, good work habits, or perseverance, as long as these traits remain constant. In a structural model, Hwang, Reed, and Hubbard (1992) showed that failing to correct for differences in worker productivity leads to significant bias in the estimation of compensating wage differentials. To reduce this bias, researchers have used panel-based approaches, such as the model estimated by Brown (1980), which is similar to:

$$W_{it} = X_{it}\beta + A_{it}\gamma + \theta_i + \varepsilon_{it},$$

where θ_i are fixed person effects that control for any persistent, static differences in wages across workers—for example, differences that might plausibly be attributable to unobserved ability or skills.

Including fixed person effects in the model changes the interpretation of the γ coefficient relative to the earlier ordinary least squares model. The fixed person effect, θ_i, absorbs the impact on wages of all the person-level characteristics, like human capital, but it also absorbs the average amenity levels of the jobs held by each person. The remaining variation in amenities not controlled by θ_i is the within-person variation over time, so this becomes the only portion of the variation in amenities used to identify γ in this approach.

What causes within-person variation in amenities? The answer can vary. Some amenities, such as the location of a job, may remain fixed over time for all jobs at a particular firm or establishment. For such amenities, the γ coefficient can only be identified by workers moving to a new job at a different location. Relying on job mobility as a source of identifying variation raises an additional set of concerns, which I discuss below, related to why workers change jobs and whether the factors that may have contributed to the job change decision are included in the wage model.

Other amenities, like whether the job is physically demanding, may also change if a worker switches occupations or tasks while remaining employed at the same firm. For these amenities, the γ coefficient on amenities can be identified by job changes across firms, within-firm changes in occupations, or a combination of both. Still other amenities, like the generosity of employer-provided health insurance benefits, may change over time even if the job and tasks performed remain the same. This type of amenity does not require any job mobility to identify γ, though job mobility may also contribute to identifying variation depending on the choice of model specification. Using a variation in amenities that is not tied to job changes can avoid one form of estimation challenge, which is related to the possibility that the error term in the wage model may influence the decision to change jobs, causing endogeneity bias.

In the panel model considered by Brown (1980), the set of amenities included whether the job requires repetitive tasks, working under stress, physical strength, working under bad physical conditions (say, extreme temperatures or physical hazards), and the risk of death on the job. The risk of death was measured at the occupation level, so changes in this risk came primarily from job changes across occupations. In the main specification, Brown finds that the compensating wage differential for the risk of death is positive and statistically significant when estimating a cross-sectional model, but the coefficient shrinks by 84 percent and becomes statistically insignificant when the same data are used to estimate a panel-based model that includes person effects.

Indeed, much (though not all[10]) of the literature that considers panel-based models finds that adding person effects to a wage model leads estimates of compensating wage differentials to shrink substantially and frequently become statistically insignificant (Kniesner et al. 2012; Viscusi and Aldy 2003). This pattern has been described as the "skills bias puzzle." The puzzle is that if job amenities are normal goods, then theory suggests that if A is a desirable amenity the bias caused by omitting θ_i is positive, and if A is an undesirable disamenity the bias is negative. However, empirical evidence has typically found the opposite.

What factors might explain the skills bias puzzle? First, the assumptions of perfect competition in the Rosen model may not hold, which will affect the estimation of compensating wage differentials, as I discuss in the next subsection. Or if workers lack information about amenity levels at the time they make job choices, then equilibrium compensating wage differentials may reflect the subjective beliefs of the marginal worker rather than objective measures of amenities. Second, in cases where some workers like A_i but others dislike it, researchers may not be able to specify a hypothesis about the direction of the marginal worker's preferences. This could potentially explain why panel-based approaches to corrected bias have unexpected impacts on estimates. One of the reasons why the literature has focused on the risk of death as an amenity is because hypotheses about the preferences of the marginal worker seem easier to justify.

When researchers are unsure about the sign of compensating wage differentials, one approach is to use data on voluntary job changes to test whether the upper and lower bounds on amenities for job changers share the same sign. Using this approach with panel data, Villanueva (2007) shows that if a worker voluntarily moves to a job with fewer amenities, the change in wages provides an upper bound on the willingness to pay for amenities. Similarly, a lower bound on preferences can be estimated among workers who voluntarily move to jobs with more amenities. For example, Villanueva (2007) finds that for workers who indicate that their job has a heavy workload, both the upper and lower bounds on γ are negative. For some other job characteristics, however, the evidence is inconclusive about the sign of compensating differentials.

[10] Two exceptions are Garen (1988) and Hwang, Reed, and Hubbard (1992).

A third explanation for the skills bias puzzle is that including person effects may not entirely control for productivity differences. If workers gain skills over time—for example, due to on-the-job training—this change in skills over time can be a source of omitted variable bias, violating the assumption that workers have conditionally homogeneous productivity. In addition to affecting wages, unmodeled changes in skills over time may be correlated with occupational promotions or job changes, which are an important source of identifying variation in amenities in some panel models.

Imperfect Competition

The early empirical work on compensating differentials, which often faced anomalous results like wrong-sided coefficients and the skills bias puzzle, prompted researchers to broaden their focus to other models of wage variation. Several studies moved away from the Rosen framework and instead combined compensating wage differentials and models of imperfect competition in labor markets by introducing features like search frictions that create wage dispersion, even among equally productive workers receiving the same amenities (Bonhomme and Jolivet 2009; Dey and Flinn 2005; Hwang, Mortensen, and Reed 1998; Lamadon, Mogstad, and Setzler 2022; Lang and Majumdar 2004; Sullivan and To 2013; Taber and Vejlin 2020).

Such studies often moved away from the types of wage regressions described above as well. One early alternative approach was to estimate preferences for amenities in imperfectly competitive labor markets by focusing on job separation rates, rather than equilibrium wages, as a dependent variable. The intuition was that how long workers remain at a job can be informative about how much a worker values the characteristics of that job (for more on this approach, see Gronberg and Reed 1994; Rosen 1986). Willingness to pay for job amenities can be estimated by dividing the effect of amenities on the probability of exiting a job by the effect of wages on the probability of exiting a job (and multiplying by –1).[11]

In an influential theoretical model by Hwang, Mortensen, and Reed (1998), workers search in a frictional labor market for jobs that offer different levels of utility from a combination of wages and nonwage amenities, where firms differ in their cost of providing amenities. Because firms that can more efficiently provide amenities at a lower cost can offer workers greater utility while still earning nonnegative profits, this leads to multiple equalizing differences functions, violating the basic structure of the Rosen equilibrium. In this scenario, even modest search frictions

[11] For example, consider an ordinary least squares regression in which the dependent variable is the probability a worker exits a job in a given year, and independent variables include the log wage and whether the job requires working in extreme conditions (say, very hot or cold temperatures). Suppose, hypothetically, the estimated coefficient on log wage is –0.8, and the coefficient on extreme job conditions is 0.2. The marginal willingness to accept a job with extreme conditions can be computed as $-1 \times 0.2/-0.8 = 0.25$, which implies that workers require 25 percent higher wages to accept a job with extreme conditions.

in the labor market cause very large bias in the estimation of workers' underlying preferences.

Imperfect competition may also help explain the skills bias puzzle. Suppose the labor market is imperfectly competitive, and three main components explain the variation in the observed data on wage and amenity pairs: (1) differences in preferences that lead workers to sort to different points along the equalizing differences function; (2) differences in worker ability that shift the equalizing differences function; and (3) heterogeneity in firm costs combined with search frictions that lead to shifts in the equalizing differences function. In this case, estimating the cross-sectional wage model described earlier will typically yield biased estimates of the γ coefficient on amenities, because the model omits the wage variation caused by the second and third components.

In addition, a panel data model with fixed worker effects will not necessarily reduce the total bias relative to the cross-sectional model. Remember that whereas the cross-sectional model uses across-worker variation caused by sorting along $M(A)$ to estimate the γ coefficient on amenities, the worker effects model uses only within-worker wage changes. In a panel data model controlling for worker effects θ_i, the variation in wages used to estimate the γ coefficient on amenities comes mainly from the third component—workers moving to jobs at firms that offer more utility because they face lower costs. However, this source of wage variation contains almost no information about the slope of the market equalizing differences function $M(A)$. When workers change jobs, they often move to jobs that offer both higher wages and more amenities. By isolating this component of the wage variation, within-worker estimates of compensating wage differentials may, in some cases, have even larger bias than cross-sectional estimates that omit worker effects (θ_i) as a control.

Matched Data Models

Rosen (1986) speculated 35 years ago that many estimation issues concerning compensating differentials would not be resolved until matched worker-firm data became available—that is, data that include linkages between each worker and each firm. Because matched employer-employee data, especially administrative censuses of jobs, contain information about how wages change when workers move between firms, they can be used to control for unobserved worker ability while also controlling for unobserved differences across firms, including unobserved job amenities or compensation policies that cause wages of similar workers to vary across firms. Matched data can also be used to relax the strong assumptions about perfect competition that are required in most cross-sectional or within-worker estimates of compensating differentials.

To gain intuition about how matched data allow estimation improvements, consider a framework that extends the panel model with worker effects so that it also includes fixed employer effects, as in:

$$W_{it} = X_{it}\beta + A_{it}\gamma + \theta_i + \psi_{J(i,t)} + \varepsilon_{it},$$

where the term $\psi_{J(i,t)}$ is a set of fixed employer effects. (The subscript $J_{(i,t)}$ can be read as employer J at which worker i is employed at time t.) The $\psi_{J(i,t)}$ term controls for static unobserved differences shared by all jobs in firm J that impact wages: for example, some firms systematically pay higher wages to reduce turnover rates. In the context of compensating differentials, the $\psi_{J(i,t)}$ term controls for compensating wage differentials associated with all of the unobserved amenities shared by jobs at firm J.[12]

This approach addresses a long-standing concern of researchers: it seems nearly impossible to quantify all of the characteristics that workers value across jobs. Including fixed employer effects reduces the need to observe every amenity, as long as the job characteristics are shared by workers at each firm and remain static over time. It is also possible to use matched data to control for unobserved factors that differ at other levels—for example $\psi_{J(i,t)}$ could be replaced with a firm-by-occupation term to control for unobserved factors that differ across occupation by firm pairs.

The conceptual goal behind this specification is to control for both unobserved worker ability and unobserved firm-level characteristics, like unobserved amenities. By including both sets of factors that can shift $M(A)$, the aim is to control for several key sources of wage dispersion among observably similar workers, so that the *conditional* equalizing differences function is analogous to the Rosen equilibrium and its slope is an estimate of compensating wage differentials. This model relaxes a large set of assumptions related to unobserved factors that drive wage variation across employers and jobs, although it still requires many assumptions about how equilibrium wages are determined.

Lavetti and Schmutte (2022) evaluate models of this form to estimate compensating wage differentials for the risk of death on the job. They show that including interactions between $\psi_{J(i,t)}$ and coarse occupation controls in the wage model appears to alleviate important sources of bias caused by unobserved differences in worker ability and unobserved employer amenities, while also relaxing the strong assumptions in the Rosen model that labor markets must be perfectly competitive. In a cross-sectional wage regression, they find that a 1 in a 1,000 increase in the annual risk of death on the job is associated with a 28 percent increase in wages. Adding worker effects to the model to control for ability causes the compensating wage differential to fall to just 3.7 percent, an 87 percent decline compared to the cross-sectional model. This pattern is consistent with the skills bias puzzle. However, using the matched structure of the data to control for firm-level differences in pay $\psi_{J(i,t)}$, the estimated compensating wage differential increases back to 17 percent. They develop a theoretical framework to show how empirical estimates from this wage model relate to the three classes of parameters in the Rosen model and propose empirical tests to assess whether parameter estimates can be interpreted as preferences or market compensating wage differentials.

[12] This model is similar to Abowd, Kramarz, and Margolis (1999), but it includes a job amenities term.

Discrimination, Exclusion, and Sorting

The Rosen model assumes a competitive market in which workers are free to access any job. Much of the literature that has relaxed the assumption of perfect competition has done so in a way that introduces nondiscriminatory frictions, like costs of job search. A smaller segment of the literature has considered a form of imperfect competition in which firms discriminate by excluding specific groups of workers from certain jobs or occupations, or by hesitating to hire them at all (Hsieh et al. 2019; Lang and Majumdar 2004). The US labor market has high levels of occupational segregation by race and gender: for example, Alonso-Villar, Del Río, and Gradín (2012) calculate the occupational Gini coefficient, which measures occupational inequality on a scale from zero (perfectly equal) to one (fully segregated), to be 0.34 by gender and 0.16 by race and ethnicity. If this occupational segregation is caused even partially by employers using demographic stereotypes to sort workers into occupations, this can change the nature of sorting in the labor market, altering the market equalizing differences function and compensating wage differentials.[13]

Note that discrimination can occur on either side of the labor market. Antos and Rosen (1975) study the case of discrimination on the labor supply side, in the case of teachers who have racially discriminatory preferences over the share of students at a school who are black or white. The authors estimate that white teachers required an additional $600 annually (in 1965 dollars) to move from an all-white school to an all-black school, while black teachers required an additional $200 annually to move in the opposite direction.

Discrimination and exclusion create some as-yet unresolved challenges for interpreting parameters from a Rosen-style model. In some social science disciplines, including social psychology, preferences are thought to be endogenously formed by social norms and stereotypes, as opposed to being exogenous primitives. From this perspective, it is difficult even to articulate conceptually what the Rosen equilibrium might look like in the absence of discrimination, including the absence of endogenous impacts on preferences (Bertrand 2020).

Quasi-random Variation in Amenities

Several recent studies have combined panel-based wage models with quasi-random variation in amenities to estimate the responsiveness of wages to amenity variation. For example, Lee and Taylor (2019) study the effects of randomized federal safety inspections of manufacturing plants carried out by the Occupational Safety and Health Administration. They find that random inspections that force plants to correct safety violations reduce plant-level fatality rates by 45 percent. In response to the safety improvement, workers' relative wages at inspected plants declined by 2–3 percent, consistent with the compensating wage differentials model.

[13] A different, though related type of scenario may occur if amenities differ across groups of workers. For example, Gruber (1994) studies the effects of mandatory changes in health insurance that increase coverage of maternity care, which leads to differential shifts in the wages of women relative to men.

In another example of quasi-random variation in amenities, Lavetti (2020) studies compensating wage differentials for the risk of death in Alaskan fisheries. This labor market displays seasonal variation in workplace hazards driven by weather patterns as well as variation from policy changes that improve safety. In addition, workers' employment contracts change frequently across fishing seasons. In this setting, it is possible to estimate compensating wage differentials while holding fixed the unobserved factors that may cause wage levels to differ across specific worker-firm pairs. In this setting, the basic cross-sectional wage model overstates compensating wage differentials by 90 percent relative to a model that accounts for unobserved job-specific heterogeneity and potentially endogenous sorting in the labor market.

Random or quasi-random variation in amenities can be useful for overcoming some of the challenges in estimating compensating wage differentials, though this type of variation does not resolve all of the issues discussed above. A benefit of random variation in amenities is that it allows researchers to estimate compensating wage differentials using changes in wages before and after the amenity variation. For example, if researchers lack rich data on worker human capital, it may still be possible to produce reliable estimates of compensating wage differentials as long as the impact of human capital on wage levels remains fixed before and after the amenity change.

In addition, random variation in amenities may change the requirements of the wage model so that, rather than requiring the error term to be uncorrelated with wage components, one only needs changes in the error term to be uncorrelated with *changes* in wage components before and after the randomization. This requirement may be easier to satisfy in some labor market settings, but it may not always hold. For example, imperfect competition that affects wage levels may also affect how responsive wages are to amenity changes, potentially biasing estimates of compensating wage differentials even with random amenity variation. Similarly, large amenity changes could lead workers to re-sort into different jobs, causing a form of endogenous variation in wage-amenity pairs in response to the random amenity variation. Despite these caveats, combining this type of amenity variation with an appropriately specified wage model is a promising direction for future progress in estimating compensating wage differentials.

Policy Applications

The Value of Statistical Life

One policy application of compensating differentials is to the "value of statistical life," which is a framework used to evaluate policies related to public safety or public health that involve the risk of death (Viscusi and Aldy 2003). Conceptually, the idea behind the value of statistical life is to imagine that many people each face a small risk of death and to determine how much money each person would be willing to pay to reduce everyone's average risk of death by a small amount such that on average one fewer person will die. The aggregate willingness to pay to prevent one expected death is the value of statistical life. For example, if 100,000 people are indifferent between accepting \$100 in additional annual earnings or a 1 in 100,000 increase in the annual risk of a fatal workplace accident, then the implied value of statistical life is \$10 million (\$100 $\times$ 100,000).

The Rosen framework has been used extensively to estimate workers' preferences for tradeoffs between wages and risks of death on the job—and the value of statistical life implied by these preferences. Many US federal regulatory agencies like the Environmental Protection Agency, the Food and Drug Administration, and the Federal Aviation Administration evaluate the costs and benefits of policies using estimates of the value of statistical life derived from labor market estimates of compensating wage differentials. For example, the US Environmental Protection Agency (2010) reviewed the literature on compensating wage differentials in labor markets and estimates of risk preferences derived from surveys where respondents were directly asked about their preferences. The agency currently places the value of saving a life at \$7.4 million (in 2006 dollars). When conducting a cost-benefit analysis of the Clean Air Act, the Environmental Protection Agency estimated that over 90 percent of the benefits came from the value of mortality reductions and were therefore based on estimates of the value of statistical life (Evans and Taylor 2020).

The labor-market setting is an appealing context to estimate the value of statistical life because it is a market with a relatively broad set of participants for whom it is common to face risk-wage tradeoffs. But based on the earlier discussion, some shortcomings of this approach become clear. Given that the value of statistical life is used to quantify society's preferences for risk reduction, consistency suggests that the value of statistical life should be calculated by aggregating workers' preferences—rather than using market compensating wage differentials. As discussed earlier, the $M(A)$ function and its slope are not a representation of workers' preferences: they are defined largely by the equilibrium sorting process that matches workers to firms.

In addition, the Rosen model identifies the preferences of the *marginal worker*. But many policies evaluated using the value of statistical life affect people who are not in the labor force (say, through safety regulations created by the Food and Drug Administration), so the risk preferences of nonworkers are not reflected in the estimates. Moreover, the marginal worker in any given job is the worker who values safety the least, which suggests that estimates of the value of statistical life derived from labor markets are likely to understate the average preferences of all workers in the labor force.

Given these limitations of the standard approaches to estimating compensating differentials, policymakers also consider a variety of alternative approaches to calculating the value of a statistical life, including using different datasets in labor market analyses as well as estimates of risk preferences derived from product markets, behaviors, or direct survey elicitation of preferences (Ashenfelter 2006; Cropper, Hammitt, and Robinson 2011).

Wage Gaps and Compensation Inequality

Earnings inequality has increased substantially in the United States and other advanced economies over the past half-century (for example, Hoffmann, Lee, and

Lemieux 2020.) Over the same time, the cost of employer-provided benefits, the largest component of which is health insurance in the United States, also increased enormously. In 2020, the worker with the median wage in the United States earned $19.15 per hour in earnings and their employers paid an additional $9.25 per hour, or 48 percent of wages, in total benefits (US Bureau of Labor Statistics 2020). These employer-provided benefits represent some of the amenities that workers might value, but additional amenities are excluded from this accounting cost, such as flexible hours and work-from-home options. Given the scale of nonwage benefits, inequality in total compensation—including the value to the worker of job amenities—may be very different than inequality in labor market earnings.

One approach to improving measures of compensation inequality would be to add compensating wage differentials to labor market earnings to calculate the market value of total compensation. For example, economists have estimated compensating wage differentials for health insurance, pension or retirement benefits, safety or workplace hazards, hours flexibility, job security, vacation or sick time, commuting time, and many other important amenities. Bell (2020) uses estimates of compensating wage differentials to construct inequality measures based on total compensation. Between black and white workers, the racial compensation gaps for workers are slightly larger than wage gaps; between males and females, the gender compensation gap may be up to two-thirds smaller than the gender wage gap. One obvious caveat to this approach is that researchers are limited to observed amenities.

In contrast to this bottom-up approach, Taber and Vejlin (2020) and Sorkin (2018) develop top-down measures of the total value of the bundle of unobserved amenities. The intuition is that if a worker voluntarily exits a job and moves to a new job that pays lower wages, the total value of amenities at that new job must exceed the value of amenities at the original job. Connecting this type of revealed preference information for all job changes using matched employer-employee data from a large set of firms, Taber and Vejlin (2020) estimate that compensating wage differentials explain a larger share of inequality in the utility of jobs than wages do. This result may still understate the full importance of compensating wage differentials because the net value of a bundle of job amenities often contains partially offsetting positive and negative wage differentials for good and bad job characteristics.

The reasons behind the dispersion of wages may affect how policymakers view these differences. If a large share of earnings inequality is caused by underlying differences in human capital, policymakers might wish to assess education and job training policies. If labor market frictions explained a large share of earnings differences, policymakers might wish to evaluate labor market institutions and regulations. However, if compensating wage differentials contribute meaningfully to earnings dispersion across groups, this variation may be less concerning for policymakers if it reflects differences in the preferences of workers who sort into jobs that differ in wages because they offer different levels of amenities that workers value.

Indeed, one could make a case that a measure of compensation that removes the effects of compensating wage differentials is *preferable* to observed income when assessing inequality. If a worker chooses to accept a lower wage and implicitly "purchases" a flexible schedule (or even a company car), should this purchase be viewed differently from an inequality perspective than the purchase of other goods or services? In this sense, netting out compensating differentials from observed earnings is consistent with a measure of earnings inequality that is agnostic regarding how people spend their total earnings. Of course, constructing such a measure is challenging and requires addressing the many estimation issues I discuss above.

Conclusion

The compensating wage differentials model is widely regarded as one of the core models of wage determination in labor economics. Empirically, compensating wage differentials are large enough to matter when economists consider topics like wage dispersion, total compensation, choices of occupations, patterns of discrimination, and even issues that may at first glance seem unrelated, like how policymakers might place a monetary value on the benefits of risk-reducing regulations. However, the literature on how to model and estimate compensating wage differentials has had to address many challenges, and has evolved considerably in recent decades.

On one side, introducing imperfect competition led the literature on compensating wage differentials to splinter as it moved away from the competitive Rosen model and considered alternative models of frictional search, bilateral bargaining, sorting, and rent-sharing. Much of this literature took a pessimistic view on whether researchers would be able to disentangle the influences of imperfect competition from compensating wage differentials.

On the other side, the expansion of matched employer-employee data, and the use of quasi-random variation in job amenities, has kindled some new optimism as researchers developed novel approaches to estimating compensating wage differentials using models that relax the assumptions of perfect competition. In particular, considering the net effect on wages of an aggregate bundle of unobserved amenities has provided new insights into the importance of the compensating wage differentials model for earnings and compensation inequality and allows researchers to consider questions that had never previously been studied. But these approaches are still nascent, and there is much to be learned.

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.

Alonso-Villar, Olga, Coral Del Río, and Carlos Gradín. 2012. "The Extent of Occupational Segregation in the United States: Differences by Race, Ethnicity, and Gender." *Industrial Relations* 51 (2): 179–212.

Antos, Joseph R., and Sherwin Rosen. 1975. "Discrimination in the Market for Public School Teachers." *Journal of Econometrics* 3 (2): 123–50.

Ashenfelter, Orley. 2006. "Measuring the Value of a Statistical Life: Problems and Prospects." *Economic Journal* 116 (510): C10–23.

Bell, Alex. 2020. "Job Amenities and Earnings Inequality." Unpublished.

Bertrand, Marianne. 2020. "Gender in the Twenty-First Century." *AEA Papers and Proceedings* 110: 1–24.

Bonhomme, Stephanie, and Gregory Jolivet. 2009. "The Pervasive Absence of Compensating Differentials." *Journal of Applied Econometrics* 24 (5): 763–95.

Brown, Charles. 1980. "Equalizing Differences in the Labor Market." *Quarterly Journal of Economics* 94 (1): 113–34.

Charles, Kerwin Kofi, Matthew S. Johnson, Melvin Stephens Jr., and Do Q. Lee. 2022. "Demand Conditions and Worker Safety: Evidence from Price Shocks in Mining." *Journal of Labor Economics* 40 (1): 47–94.

Cropper, Maureen, James K. Hammitt, and Lisa A. Robinson. 2011. "Valuing Mortality Risk Reductions: Progress and Challenges." *Annual Review of Resource Economics* 3 (1): 313–36.

DeLeire, Thomas, Shakeeb Khan, and Christopher Timmins. 2013. "Roy Model Sorting and Nonrandom Selection in the Valuation of a Statistical Life." *International Economic Review* 54 (1): 279–306.

Dey, Matthew S., and Christopher J. Flinn. 2005. "An Equilibrium Model of Health Insurance Provision and Wage Determination." *Econometrica* 73 (2): 571–627.

Duncan, Greg J. 1976. "Earnings Functions and Nonpecuniary Benefits." *Journal of Human Resources* 11 (4): 462–83.

Ekeland, Ivar, James J. Heckman, and Lars Nesheim. 2004. "Identification and Estimation of Hedonic Models." *Journal of Political Economy* 112 (S1): S60–109.

Evans, Mary F., and Laura O. Taylor. 2020. "Using Revealed Preference Methods to Estimate the Value of Reduced Mortality Risk: Best Practice Recommendations for the Hedonic Wage Model." *Review of Environmental Economics and Policy* 14 (2): 282–301.

Flood, Sarah, Miriam King, R. Rodgers, Steven Ruggles, R. Warren, and M. Westberry. 2022. "Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset]." IPUMS. https://doi.org/10.18128/D030.V10.0 (accessed December 14, 2022).

Garen, John. 1988. "Compensating Wage Differentials and the Endogeneity of Job Riskiness." *Review of Economics and Statistics* 70 (1): 9–16.

Gronberg, Timothy J., and W. Robert Reed. 1994. "Estimating Workers' Marginal Willingness to Pay for Job Attributes Using Duration Data." *Journal of Human Resources* 29 (3): 911–31.

Gruber, Jonathan. 1994. "The Incidence of Mandated Maternity Benefits." *American Economic Review* 84 (3): 622–41.

Hersch, Joni. 2011. "Compensating Differentials for Sexual Harassment." *American Economic Review* 101 (3): 630–34.

Hoffmann, Florian, David S. Lee, and Thomas Lemieux. 2020. "Growing Income Inequality in the United States and Other Advanced Economies." *Journal of Economic Perspectives* 34 (4): 52–78.

Hotz, V. Joseph, Per Johansson, and Arizo Karimi. 2018. "Parenthood, Family Friendly Firms, and the Gender Gaps in Early Work Careers." NBER Working Paper 24173.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow. 2019. "The Allocation of Talent and U.S. Economic Growth." *Econometrica* 87 (5): 1439–74.

Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed. 1998. "Hedonic Wages and Labor Market Search." *Journal of Labor Economics* 16 (4): 815–47.

Hwang, Hae-shin, W. Robert Reed, and Carlton Hubbard. 1992. "Compensating Wage Differentials and Unobserved Productivity." *Journal of Political Economy* 100 (4): 835–58.

Kniesner, Thomas J., W. Kip Viscusi, Christopher Woock, and James P. Ziliak. 2012. "The Value of Statistical Life: Evidence from Panel Data." *Review of Economics and Statistics* 94 (1): 74–87.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market." *American Economic Review* 112 (1): 169–212.

Lang, Kevin, and Sumon Majumdar. 2004. "The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications." *International Economic Review* 45 (4): 1111–28.

Lavetti, Kurt. 2020. "The Estimation of Compensating Wage Differentials: Lessons from the Deadliest Catch." *Journal of Business and Economic Statistics* 38 (1): 165–82.

Lavetti, Kurt, and Ian M. Schmutte. 2022. "Estimating Compensating Wage Differentials with Endogenous

Job Mobility." Unpublished.
Lee, Jonathan M., and Laura O. Taylor. 2019. "Randomized Safety Inspections and Risk Exposure on the Job: Quasi-experimental Estimates of the Value of a Statistical Life." *American Economic Journal: Economic Policy* 11 (4): 350–74.
Lucas, Robert E. B. 1977. "Hedonic Wage Equations and Psychic Wages in the Returns to Schooling." *American Economic Review* 67 (4): 549–58.
Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements." *American Economic Review* 107 (12): 3722–59.
Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1): 34–55.
Rosen, Sherwin. 1986. "The Theory of Equalizing Differences." In *Handbook of Labor Economics*, Vol. 1, edited by Orley C. Ashenfelter and Richard Layard, 641–92. Amsterdam: North Holland.
Smith, Adam. 1776. *The Wealth of Nations.* New York: Modern Library, 2000.
Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *Quarterly Journal of Economics* 133 (3): 1331–93.
Sullivan, Paul, and Ted To. 2013. "Job Dispersion and Compensating Wage Differentials." US Bureau of Labor Statistics Working Paper 469.
Taber, Christopher, and Rune Vejlin. 2020. "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market." *Econometrica* 88 (3): 1031–69.
Thaler, Richard, and Sherwin Rosen. 1976. "The Value of Saving a Life: Evidence from the Labor Market." In *Household Production and Consumption*, edited by Nestor E. Terleckyj, 265–302. New York: NBER.
US Bureau of Labor Statistics. 2020. "Employer Costs for Employee Compensation—March 2020." https://www.bls.gov/news.release/archives/ecec_06182020.pdf (accessed March 7, 2023).
US Environmental Protection Agency. 2010. "Valuing Mortality Risk Reductions for Environmental Policy: A White Paper." https://www.epa.gov/environmental-economics/valuing-mortality-risk-reductions-environmental-policy-white-paper-2010.
Villanueva, Ernesto. 2007. "Estimating Compensating Wage Differentials Using Voluntary Job Changes: Evidence from Germany." *Industrial and Labor Relations Review* 60 (4): 544–61.
Viscusi, W. Kip, and Joseph E. Aldy. 2003. "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World." *Journal of Risk and Uncertainty* 27 (1): 5–76.