Pricing Job Amenities: A Practitioner's Manual

Alex Bell

Department of Economics, Georgia State University

ambell@gsu.edu

Abstract

This paper provides a practical framework for estimating compensating differentials from observational data. Observed wage–amenity relationships combine two sources of variation: differences in the opportunities workers receive, traditionally treated as nuisance variation that confounds estimation, and differences in how compensation is divided between wages and amenities, which in principle identify the compensating differential. Conventional instrumental variables methods seek an observed variable related to the identifying variation. This paper shows that identification can instead proceed from an observed variable related to the nuisance variation—an anti-instrumental variable. This inversion is consequential in practice: measures related to workers' opportunities are widely available, while amenity-specific instruments are rare. The economic screen for an anti-instrument closely parallels how researchers already select controls for nuisance variation in hedonic applications, but does not require observing every source of opportunity differences. The paper offers a practical roadmap using the `aivreg` Stata package. A companion context pack, available from the author upon request, presents the roadmap in a more granular, AI-readable format. Applying the method to the replication data from Mas and Pallais (2017) reverses the counterintuitive signs in conventional hedonic estimates: across the three focal amenities, the AIV prices have the signs implied by experimentally elicited preferences and fall within the corresponding willingness-to-pay distributions.

Keywords: compensating differentials; job amenities; labor markets; hedonic wage models; anti-instrumental variables

JEL Classification: J31, J24, C21, J81

1 Introduction

Job amenities are part of compensation. Yet estimating their prices from observational data remains difficult because observed wage–amenity relationships combine two economically distinct sources of variation. Several strands of literature support the premise that workers differ in the opportunities governing the compensation bundles they can attain (Mincer, 1974; Burdett and Mortensen, 1998; Hwang et al., 1998; Cunha and Heckman, 2007; Chetty and Hendren, 2018). Compensation may also be divided differently between wages and amenities. This second source of variation—compensation composition—traces the compensating differential. The first confounds it: workers with better opportunities may obtain both higher wages and more desirable amenities, concealing or even reversing the underlying trade-off. The resulting sorting problem is central to modern work on compensating differentials and firm value (Sorkin, 2018; Lavetti, 2023; Mas, 2025).

This distinction mirrors the familiar problem of supply and demand. When the slope of demand is the object of interest, demand shocks are nuisance variation: they confound estimation of the demand curve. Supply shocks instead move equilibrium along the demand curve and provide identifying variation. The instrumental-variables method uses an observed variable related to those supply shocks to recover the demand curve. If supply were the object of interest, the roles would reverse. Variation is therefore not inherently nuisance or identifying; its role depends on the equation being estimated.

The compensating-differentials problem has the same structure. Differences in opportunities shift the observed wage–amenity relationship, while differences in compensation composition trace the relationship of interest. An ordinary instrumental-variables strategy would therefore seek a variable that shifts an amenity without otherwise changing workers' opportunities. Randomly assigned OSHA inspections, for example, have been used to instrument plant fatality risk (Lee and Taylor, 2019). Yet credible experiments or quasi-experiments that induce exogenous changes in a particular amenity or its cost remain rare (Mas, 2025). Standard labor-market datasets, by contrast, often contain schooling, test scores, and employment histories—variables that can be screened as measures of accumulated opportunity (Farber and Gibbons, 1996; Altonji and Pierret, 2001; Lubinski et al., 2014; Deming et al., 2016; Sorkin, 2018; Bell et al., 2019; Chetty et al., 2020). Their availability does not by itself establish AIV validity. This paper shows how to use information about that nuisance variation for identification rather than only attempt to control it away. I call such a variable an anti-instrumental variable.

In a conventional hedonic regression, the variation researchers hope remains after con-

trols is the compensation-composition variation that identifies the price. A variable that independently shifts this margin would therefore absorb identifying variation rather than remove nuisance variation. The economic screen for an anti-instrument is closely related to this familiar logic for selecting controls. A candidate must contain information about differences in workers' opportunities. But once the relevant opportunity component is held fixed, it should contain no additional information about how compensation is divided between wages and amenities. The important departure from a conventional control strategy is that the researcher need not observe enough variables to eliminate all remaining opportunity variation. In the linear model, one valid but incomplete measure of opportunities can identify the compensating differential.

This paper turns that idea into a practical roadmap. It explains how to select and evaluate an anti-instrument, estimate amenity prices with the companion `aivreg` Stata package, and interpret the resulting estimates. A companion context pack, available from the author upon request, presents the roadmap in a more granular, AI-readable format. The formal identification results underlying the approach are developed in Bell (2020) and Bell et al. (2024); the narrower purpose here is to make the method accessible to researchers studying non-wage amenities.

The paper also asks whether the theoretical solution works in practice. I apply the method to the replication data from Mas and Pallais (2017), whose experiment provides an external benchmark for the observational price estimates. Using only the observational data, the AIV prices have the signs implied by experimentally elicited preferences and fall within the corresponding willingness-to-pay distributions. The application therefore illustrates both how to implement the method and how experimental evidence can be used to assess whether it recovers economically meaningful compensating differentials.

2 Estimation

Suppose that workers choose among jobs offering wage w and amenity z. Observed compensation bundles differ for two conceptually distinct reasons: workers face different opportunities, and compensation is divided differently between wages and amenities.[1] The researcher wants to estimate the signed equilibrium price β: the wage difference associated with the amenity among jobs reached after comparable opportunity realizations. For a desirable amenity, β will ordinarily be negative.

[1] This section considers a single amenity for ease of exposition, but the estimator extends directly to multiple amenities.

The two sources of variation can be represented symmetrically as

$$w_i = \alpha_\eta + \beta z_i + r_{\eta i}, \tag{1}$$

$$w_i = \alpha_\theta + \Gamma z_i + q_{\theta i}. \tag{2}$$

The first equation defines the compensating differential of interest. Its residual $r_{\eta i}$ summarizes the net realization of the potentially many forces that shape worker i's opportunities. It is a residual coordinate for this equation, not a claim that opportunity is one-dimensional or that amenities causally determine wages.[2] The second equation records the other equilibrium direction through compensation space. Its residual $q_{\theta i}$ indexes compensation-composition variation, and Γ is not the amenity price.[3]

At a given value of r_η, compensation-composition variation identifies β. The standard regression of w on z is confounded when z is also related to r_η. Workers with better opportunities may obtain both higher wages and more desirable amenities, concealing or even reversing the underlying trade-off. This includes the familiar problem of ability bias, but is broader: education, skills, search frictions, luck, discrimination, firms, matches, and other stochastic events may all affect the compensation bundles workers can attain. Indeed, Mas (2025) concludes that omitted productivity measures are "a primary source of bias in compensating differential estimates."

The researcher does not observe r_η directly, but may observe a variable h containing information about workers' opportunities. The estimator can be represented using the following two-equation procedure:

$$h_i = \pi_0 + \pi_w w_i + \pi_z z_i + \nu_i, \tag{3}$$

$$w_i = \alpha + \beta z_i + \gamma \hat{h}_i, \tag{4}$$

where $\hat{h}_i$ is the fitted value from Equation 3. Under the assumptions described in the next section, $\hat{h}_i$ reconstructs an affine transformation of the opportunity residual $r_{\eta i}$, which is sufficient to account for it in the second equation. Equivalently, the price can

[2]Earlier versions of this paper imposed a single-index restriction under which the observed wage–amenity bundle determined a scalar measure of latent job quality. The present formulation does not require that restriction. Opportunity and compensation-composition shocks may each be multidimensional; r_η is simply the residual coordinate of the maintained price equation. The earlier restriction was therefore stronger than necessary for identification.

[3]The analogous inverse-demand and inverse-supply system is $p = \alpha_D + \beta_D q + \varepsilon_D$ and $p = \alpha_S + \beta_S q + \varepsilon_S$. When the demand curve is the object of interest, demand shocks ε_D confound its estimation, whereas supply shocks ε_S move equilibrium along it and identify β_D. In the wage–amenity system, r_η plays the role of the nuisance residual in the target equation, while q_θ indexes the compensation-composition direction that identifies β.

be recovered directly from the first equation as

$$\beta = -\frac{\pi_z}{\pi_w}.$$

Equation 4 is therefore an alternative representation of the same coefficient-ratio estimator rather than a separate source of identification.

To build intuition, consider the restrictive special case in which h equals r_η plus an independent, mean-zero disturbance. Noise in the dependent variable does not change the population coefficients in Equation 3. Because $r_\eta = w - \beta z$ after absorbing the intercept, the fitted values reconstruct the opportunity residual, and normalizing by the wage coefficient recovers β.[4]

Because the estimator is a ratio of coefficients, weak information about the opportunity residual creates an inference problem analogous to weak instruments. The freely available `aivreg` Stata package reports the coefficient-ratio estimate and provides Anderson–Rubin confidence sets robust to weak identification. With multiple amenities, Equation 3 includes the full observed amenity vector, and each amenity price is obtained by normalizing its coefficient by the wage coefficient.

The same geometric insight can extend beyond an affine relationship between h and r_η. If the conditional mean of h is an injective function of $r_\eta = w - \beta z$, its contours identify compensation bundles with the same opportunity residual, and their wage-normalized slope identifies β. The formal conditions and proof for flexible identification are developed in Bell (2020).

Like other hedonic methods using observational data, AIV still requires the researcher to observe the amenities being priced and to specify a coherent price over the relevant support; it cannot price amenities that are not observed. Its practical advantage is narrower but consequential: the researcher need not observe the complete stochastic history determining workers' opportunities. One valid but incomplete opportunity variable can be sufficient under the assumptions stated next.

Functional form presents an additional limitation. The linear coefficient-ratio estimator combines a constant-price equation with an affine relationship between the candidate opportunity variable and r_η. The contour argument above can relax the latter restriction, but does not by itself resolve heterogeneity in the amenity price. When prices vary across compensation space, a global linear coefficient may combine local slopes in ways that depend on sorting and need not have a simple average-price interpretation.

[4]Within this special case, including noisy h directly as a control leaves residual opportunity confounding. Greater noise makes that conventional control strategy less effective. For AIV, under the maintained assumptions, weaker information in h instead produces less precise identification of β. The general method does not require h to be a literal noisy measurement of r_η.

Local and flexible estimation offer a possible route forward, but require adequate support and less turnkey methods of estimation and inference.

3 Assumptions

Before turning to the assumptions specific to AIV, it is useful to separate several requirements shared with other hedonic methods. The amenity price must be a coherent economic object over the relevant support; the researcher must observe the amenity bundle being priced; and the chosen functional form must adequately represent the price on that support. AIV does not remove these requirements. Its distinctive contribution concerns the information needed to separate differences in workers' opportunities from differences in the composition of compensation.

The key identifying assumption is an AIV exclusion restriction. In the one-amenity notation of Section 2, it can be stated as

$$h \perp q_\theta \mid r_\eta.$$

Once the opportunity residual is held fixed, the candidate AIV contains no independent information about how compensation is divided between wages and amenities.[5] This is the AIV counterpart to IV exogeneity: after holding fixed the nuisance residual that h is intended to reveal, h supplies no additional information about the identifying variation.[6] The restriction does not require h and the amenity to be unrelated unconditionally. A candidate may predict amenities through its relationship with workers' opportunities; it fails the exclusion restriction only if it supplies an additional compensation-composition direction once r_η is fixed.[7]

The second requirement is relevance, analogous to the relevance condition for an ordinary instrument. The candidate h must be informative about r_η, whereas an ordinary instrument must instead be informative about q_θ. For flexible identification, its conditional mean must vary sufficiently systematically across values of r_η to distinguish constant-opportunity contours. For the linear coefficient-ratio estimator, that relationship is assumed to be affine with a nonzero slope. Weak relevance primarily appears as

[5] With multiple amenities, let R_θ collect the remaining compensation-composition directions. The corresponding condition is $h \perp R_\theta \mid r_\eta$.

[6] The analogous IV condition requires a variable to supply information about the identifying variation but not the nuisance variation. In this sense, IV and AIV reverse the roles assigned to the two sources of variation.

[7] Full conditional independence is stronger than necessary. Flexible identification requires the conditional mean of h to remain unchanged across compensation-composition directions once r_η is fixed, together with sufficient variation in that conditional mean across values of r_η. In the linearized model, these conditions reduce to the corresponding conditional-covariance restriction and affine relevance.

imprecise or uninformative confidence sets under the maintained model rather than as a reason to treat a weak candidate as a more complete control.

Neither requirement imposes a primitive one-dimensional model of workers' opportunities. The scalar r_η is the residual coordinate of the maintained price equation and may summarize a multidimensional stochastic history involving education, skills, health, search frictions, discrimination, firms, matches, luck, and other forces. Correspondingly, a candidate AIV may be an accumulated realization of earlier opportunities, a measure of the current opportunity state, a cause of later opportunities, or some combination of these. Premarket timing is neither required nor sufficient. The researcher must instead defend relevance and the exclusion restriction for the particular candidate, amenity bundle, population, and setting.

When several candidate AIVs are available, overidentification tests can provide compatibility checks among the maintained restrictions. Rejection indicates that the maintained restrictions cannot all hold simultaneously, but non-rejection does not establish the exclusion restriction; economic reasoning remains necessary. The same logic can extend beyond job amenities when the researcher can define a coherent target equation, distinguish its nuisance residual from the remaining identifying directions, and identify a candidate satisfying the corresponding relevance and exclusion requirements.

4 Application

To illustrate the estimator in practice, I revisit Mas and Pallais (2017), a landmark study on alternative work arrangements that has shaped how economists think about non-wage amenities. Their paper overcame what many see as the limitations of observational cross-sectional data by partnering with an employer to give real job candidates choices between different packages of earnings and amenities. It remains among the most frequently cited papers in the literature.[8]

Many readers will view the type of hedonic regressions in Table 1 of Mas and Pallais (2017) as a canonical approach: a regression of wages on amenities with extensive controls for worker characteristics. This short application first asks whether applying the anti-IV estimator changes the counterintuitive signs of those results, and then whether the resulting equilibrium price estimates have the signs and locations implied by the willingness-to-pay distributions elicited in the experiment.

Mas and Pallais (2017) noted that their hedonic regressions produced counterintuitive

[8]For instance, it is the most frequently cited paper in the recent *Handbook of Labor Economics* chapter on non-wage amenities.

Table 1: Conventional and AIV Price Estimates in the Mas and Pallais (2017) Setting

	Separate Regressions		Same Regression	
	Conventional	AIV	Conventional	AIV
Schedule flexibility				
Can vary times at which workday starts or ends	0.0633	-0.564	0.0604	-0.512
	(0.00983)	(0.0345)	(0.00933)	(0.0333)
Work from home				
Does any work from home	0.0796	-1.651		
	(0.0140)	(0.0753)		
Formal work from home arrangement	0.100	-0.807	0.0708	-0.583
	(0.0264)	(0.0724)	(0.0267)	(0.0748)
Irregular schedule				
Works an irregular schedule	-0.0702	0.341	-0.0660	0.307
	(0.0114)	(0.0414)	(0.0114)	(0.0400)
Works an irregular but consistent schedule	-0.0525	0.518		
	(0.0124)	(0.0517)		
Works an irregular, inconsistent schedule	-0.0795	-0.00526		
	(0.0192)	(0.0504)		

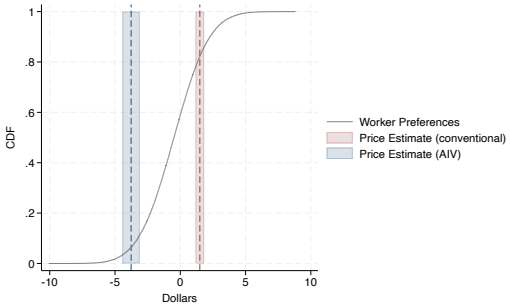
The first column of the table exactly reproduces the first column of Table 1 in Mas and Pallais (2017). It shows the results of separate regressions of log weekly earnings on the given amenity after controlling for: "hours worked per week in a respondent's main job, an indicator for working part-time, race/ethnicity, gender, educational attainment, marital status, geographic region, age, self-employment, and an indicator for being born outside the United States." The second column also relates log earnings to each amenity in separate regressions in the same dataset, but uses the CPS educational-attainment recode as an anti-instrument rather than as a control and includes no additional controls. This measure records schooling and degree level rather than field of study. The final two columns repeat this contrasting analysis when pricing a parsimonious set of amenities (one from each category) jointly rather than in separate regressions. Standard errors are robust and at the individual level. AIV standard errors are calculated using Anderson-Rubin finite-sample-robust inference performed by the `aivreg` Stata command.

results: more pleasant work arrangements were associated with higher wages, contrary to the notion of compensating differentials. They write that "Table 1 shows the difficulty of estimating compensating differentials for the work arrangements we study." Despite controlling "for a variety of worker (and some job) characteristics" they continue to find "more pleasant work arrangements are correlated with higher wages." The unexpected signs reflect the identification problem developed above: observed wage–amenity relationships combine differences in workers' opportunities with differences in how compensation is divided between wages and amenities. An extensive but incomplete control set need not remove the remaining opportunity residual.

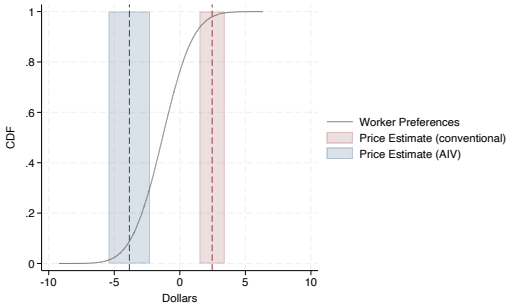
Table 1 presents the comparison. Column (1) reproduces the results presented in Mas and Pallais (2017, Table 1, Column 1), which uses educational attainment and other observed worker characteristics as conventional controls. The results imply that workers must be paid more to accept a flexible schedule or to work from home, while irregular schedules appear to be amenities for which workers accept lower pay. As the authors note, these conclusions run opposite to common sense and also seem unlikely to be correct given the preference data collected in the same paper. Column (2) instead uses the CPS educational-attainment measure as a candidate AIV. The measure may summarize accumulated earlier opportunities, cause or mediate later opportunities, or do both; the application does not take a stand among these interpretations. The maintained assumption is that it is informative about the opportunity residual but contains no additional information about how compensation is divided once that residual is held fixed. In five of six cases, the signs flip: flexibility and work from home are priced as amenities that workers must sacrifice wages to obtain, while irregular schedules carry a wage premium. Pricing one amenity from each category jointly changes magnitudes but not signs. Throughout, the conventional approach and the anti-IV approach systematically deliver opposite conclusions.

The sign reversals are the most robust feature of the exercise. Exact magnitudes warrant more caution: they vary across amenity definitions and between separate and joint specifications; the linear estimator summarizes prices that may vary at different levels of overall opportunity; and translating experimental willingness-to-pay evidence into an equilibrium-price benchmark requires distributional and market-incidence assumptions whose uncertainty is not included in the plotted confidence intervals. The exercise therefore places greater weight on the signs and broad locations of the estimates within the willingness-to-pay distributions than on exact level agreement with a single benchmark.

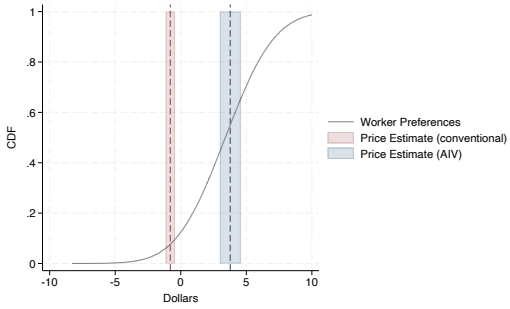
Figure 1: Preferences and Prices for Amenities



(a) Flexible Schedule



(b) Work from Home



(c) Irregular Schedule

Notes: Curves in black are the distributions of WTP reported in Table 5 of Mas and Pallais (2017) for the following amenities: Flexible Schedule, Work from Home, and Employer Discretion. The conventional pricing models shown by red lines (and shaded 95% confidence intervals) represent the price estimate obtained by regressing hourly earnings on the three amenities with the controls used by Mas and Pallais (2017). This specification differs from that of Column 3 of Table 1 in that the outcome of log weekly earnings has been replaced by hourly earnings; the change of outcome is necessary because Mas and Pallais (2017) reported willingness to pay distributions in dollars per hour as opposed to in percentages. The AIV price estimates shown in blue use the CPS educational-attainment measure as the anti-instrument when relating hourly earnings to the given amenity, analogous to Column 4 of Table 1.

The equilibrium price is not generally expected to equal average or median willingness to pay. In the standard binary-amenity model, it reflects the valuation of the marginal worker—the worker just indifferent between jobs with and without the amenity. When a desirable amenity is rare, workers must value it unusually highly to accept the wage sacrifice required to obtain it. The marginal worker's willingness to pay, and therefore the equilibrium price, can consequently lie far from the average without contradicting the experimental evidence.

Figure 1 relates these last two columns of price estimates directly to the preference data elicited by Mas and Pallais (2017). Panel A shows that most workers positively value flexible schedules, yet the conventional model suggests wages are higher for those with flexible jobs. In contrast, the AIV estimate prices flexibility as an amenity for which workers must sacrifice wages. Because flexible jobs remained uncommon despite strong worker demand, the equilibrium evidence points to substantial employer-side costs or constraints on provision. Panel B reports a similar pattern for working from home. Workers report substantial willingness to pay for remote work, and the AIV estimate shows that they must give up wages to obtain it. Because formal work from home was rare, this marginal-worker logic predicts a substantial wage sacrifice even before considering the exact magnitude of the estimate. Under the hedonic framework, this combination implies that employer-side costs or constraints, rather than weak worker demand, limited remote work at the time. The conventional model, again, gives the wrong sign.

Panel C considers irregular schedules. Workers overwhelmingly prefer to avoid jobs with irregular hours, yet the conventional model implies they are willing to accept lower wages in such jobs. The AIV estimates instead price irregular schedules with a wage premium, consistent with the standard Rosen (1974) hedonic framework in which equilibrium amenity prices reflect both worker preferences and employer production technologies. Some employers (hospitals, police forces) may require night and weekend work as part of their technology, while others (office-based employers) produce more efficiently with regular schedules. The AIV estimates fall in the middle of the preference distribution, consistent with this heterogeneity in employer costs.

In summary, prior work was correct to highlight that differences in workers' opportunities confound conventional hedonic regressions and that experimental evidence can provide credible benchmarks for amenity valuation. But opportunity confounding need not be a limitation inherent to cross-sectional observational data. Using the same data, AIV treats the educational-attainment measure not as an incomplete control but as information about the opportunity residual. Across the three focal amenities, the resulting equilibrium price estimates have the signs implied by experimentally elicited

preferences and fall within the corresponding willingness-to-pay distributions. More broadly, these results reinforce the central contribution of Mas and Pallais (2017): experimental benchmarks are invaluable for understanding job amenities, and methods that can recover those patterns in non-experimental settings substantially expand their applicability.

5 Conclusion

A growing number of applied papers have used AIV to price job amenities in a wide array of countries and contexts (Bell, 2020; Folke and Rickne, 2022; Burbano et al., 2023; Ferreira et al., 2025; Bell et al., 2023; Maideu-Morera, 2024; De Schouwer et al., 2025; De Schouwer and Kesternich, 2025; Darougheh et al., 2025). This paper provides a practical account of why the method works and how to use it. Observed wage–amenity bundles combine differences in workers' opportunities with differences in how compensation is divided. Conventional controls attempt to remove the former; AIV instead uses a valid, possibly incomplete measure of opportunities to reveal the nuisance residual. This change in how the same information is used substantially reduces the information burden for estimating compensating differentials from observational data.

The comparison reveals a broader duality between conventional controls and AIV. Both approaches favor variables that contain substantial information about workers' opportunities but little independent information about compensation composition. Weak opportunity information leaves conventional estimates biased but makes AIV estimates imprecise. Information about compensation composition absorbs identifying variation and reduces precision when used as a conventional control, but violates the AIV exclusion restriction and biases the AIV estimate. The two approaches therefore assign the same information opposite econometric roles. Most importantly, AIV recasts what prior work identifies as a primary source of bias in conventional hedonic estimates—incomplete measurement of workers' opportunities—as weak identification that is visible in the resulting confidence set, provided the exclusion restriction holds.

These results also open several questions for future research. Within amenity pricing, priorities include flexible estimation when prices vary across levels of overall opportunity; principled selection or grouping when measured amenities are close substitutes or otherwise highly correlated; and sensitivity analysis showing how small violations of the AIV exclusion restriction translate into price bias. Additional comparisons with experimental and quasi-experimental benchmarks can clarify when local evidence on willingness to pay supports inference about equilibrium prices. Beyond job amenities, the same residual duality may prove useful whenever observational outcomes combine nui-

sance variation with the variation that identifies an economic relationship. These are questions made possible by the method, not caveats that undo it. Under its stated assumptions, AIV makes it possible to recover equilibrium amenity prices from commonly available observational data without waiting for rare amenity-specific instruments.

Most of all, the method reframes how observational and experimental evidence fit together. Elicited willingness-to-pay distributions describe the demand side of the hedonic market; a credibly identified equilibrium price locates where that demand meets the costs of provision. Such a price therefore completes the experimental willingness-to-pay agenda rather than competing with it: placed beside a willingness-to-pay distribution, it closes the hedonic system and brings into view the employer-side costs and constraints that neither preference evidence nor conventional wage regressions reveal on their own. Job amenities can then be studied not only through what workers would pay, but through what the market charges.

References

Altonji, Joseph G. and Charles R. Pierret, "Employer Learning and Statistical Discrimination," *The Quarterly Journal of Economics*, 2001, *116* (1), 313–350.

Bell, Alex, "Job Amenities & Earnings Inequality," *MPRA Paper*, May 2020. Number: 118516.

_ **, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen**, "Who Becomes an Inventor in America? The Importance of Exposure to Innovation," *The Quarterly Journal of Economics*, 2019, *134* (2), 647–713.

_ **, Sophie Calder-Wang, and Shusheng Zhong**, "Pricing Neighborhood Amenities: A Proxy-Based Approach," September 2023.

_ **, Stephen B. Billings, Sophie Calder-Wang, and Shusheng Zhong**, "An Anti-IV Approach for Pricing Residential Amenities: Applications to Flood Risk *," July 2024.

Burbano, Vanessa C., Olle Folke, Stephan Meier, and Johanna Rickne, "The Gender Gap in Meaningful Work," *Management Science*, December 2023.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–273.

Chetty, Raj and Nathaniel Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *The Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

_ , **John N. Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan**, "Income Segregation and Intergenerational Mobility across Colleges in the United States," *The Quarterly Journal of Economics*, 2020, *135* (3), 1567–1633.

Cunha, Flavio and James J. Heckman, "The Technology of Skill Formation," *American Economic Review*, 2007, *97* (2), 31–47.

Darougheh, Saman, Alexander Dietrich, and Jonna Olsson, "Workplace Amenities and the Gender Wage Gap: Evidence from Denmark," September 2025.

Deming, David J., Noam Yuchtman, Amira Abulafi, Claudia Goldin, and Lawrence F. Katz, "The Value of Postsecondary Credentials in the Labor Market: An Experimental Study," *American Economic Review*, 2016, *106* (3), 778–806.

Farber, Henry S. and Robert Gibbons, "Learning and Wage Dynamics," *The Quarterly Journal of Economics*, 1996, *111* (4), 1007–1047.

Ferreira, Susana, Sara Martínez de Morentin, and Amaya Erro-Garcés, "Measuring job risks when hedonic wage models do not do the job," *Journal of Environmental Economics and Management*, January 2025, p. 103120.

Folke, Olle and Johanna Rickne, "Sexual Harassment and Gender Inequality in the Labor Market*," *The Quarterly Journal of Economics*, November 2022, *137* (4), 2163–2212.

Hwang, Haeshin, Dale T. Mortensen, and W. Robert Reed, "Hedonic Wages and Labor Market Search," *Journal of Labor Economics*, October 1998, *16* (4), 815–847.

Lavetti, Kurt, "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications," *Journal of Economic Perspectives*, August 2023, *37* (3), 189–212.

Lee, Jonathan M. and Laura O. Taylor, "Randomized Safety Inspections and Risk Exposure on the Job: Quasi-experimental Estimates of the Value of a Statistical Life," *American Economic Journal: Economic Policy*, 2019, *11* (4), 350–374.

Lubinski, David, Camilla P. Benbow, and Harrison J. Kell, "Life Paths and Accomplishments of Mathematically Precocious Males and Females Four Decades Later," *Psychological Science*, 2014, *25* (12), 2217–2232.

Maideu-Morera, Gerard, "Amenity-Biased Technical Change," 2024.

Mas, Alexandre, "Non-Wage Amenities," April 2025.

_ **and Amanda Pallais**, "Valuing Alternative Work Arrangements," *American Economic Review*, December 2017, *107* (12), 3722–3759.

Mincer, Jacob A., *Schooling, Experience, and Earnings*, National Bureau of Economic Research, 1974.

Schouwer, Thimo De and Iris Kesternich, "Work Meaning and the Flexibility Puzzle," *Journal of Labor Economics (Forthcoming)*, 2025.

_ **, Thibault Deneus, and Marco Forti**, "How to Make Work Meaningful?," April 2025.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, August 2018, *133* (3), 1331–1393.