Pricing Job Amenities: A Practitioner's Manual

By Alex Bell
Department of Economics, Georgia State University
ambell@gsu.edu

This paper provides a practical framework for estimating compensating differentials in the presence of unobserved worker ability. The approach treats ability as a structured residual that shapes access to both wages and non-pay job attributes. Job choice is assumed to be governed by a single vertical job-quality index, and—conditional on this index—an observed proxy for ability acts as an "anti-instrument" that allows researchers to recover amenity prices from standard observational data. Identification relies on a simple conditional independence assumption: if pay, amenities, and underlying job quality were fully observed, the anti-instrument would add no additional information. The paper emphasizes implementation and interpretation rather than technical proofs, offering a practical roadmap for applied researchers; adoption is supported by the companion `aivreg` *Stata package, which enables estimation using standard cross-sectional datasets. An application to the experimental setting of Mas and Pallais (2017) shows that the method reproduces experimentally measured willingness-to-pay patterns, illustrating how observational data can recover credible estimates of job amenity valuations when differences in workers' opportunity sets are properly accounted for.*

I. Introduction

When economists estimate production functions, it is now standard to treat productivity as a structured residual—something unobserved, but correlated with inputs. This age-old insight, attributed to Solow (1956), transformed empirical work in macroeconomics and industrial organization. Labor economics faces an analogous problem: wages and job characteristics ("amenities") are both determined by unobserved ability. But just as productivity can be treated as a residual in production functions, ability can be treated as a residual in wage-amenity models. Yet many papers still regress wages on amenities with controls, as if ability were either observed or ignorable in relation to the amount of amenities a worker can obtain.

This paper presents a simple solution. Building on a latent index framework, I show how a noisy proxy for ability can be used to estimate the wage-amenity

trade-off. The key assumption is that ability enters selection in a single-index way—a standard assumption in both structural and reduced-form work. Unlike structural models, however, this approach avoids functional form assumptions. And unlike typical reduced-form strategies, it does not rely on natural experiments. Like typical hedonic methods, it is appropriate when the researcher believes they can observe the set of utility-generating attributes that workers choose between.

The estimator is introduced and formally justified in a series of longer technical papers. Bell (2020) includes proofs and applications to gender, race, and class pay gaps, while Bell et al. (2024) extends the model to residential amenities and includes additional technical components. This paper is intended as a concise practical companion to, rather than a substitute for, the formal theory and empirical applications developed in those papers. The goal here is narrower: to offer a practitioner's guide for those who want to do research on non-wage amenities. Researchers can estimate the model with standard software via the companion `aivreg` package.

II. Estimation

Suppose that workers choose jobs based on a willingness to pay for an amenity and a budget constraint that varies across people.[1] With a slight loss of generality, we can refer to the budget constraint as "ability."[2] Suppose the researcher wants to quantify the trade-off workers face when choosing between wage w and a particular amenity z, and the researcher observes these variables. The researcher cannot directly observe ability η, but observes some empirical estimate of it, h.

The standard regression of w on z is biased because η is an omitted variable that causes both the outcome and regressor. In other words, higher-ability individuals can obtain both more wages and amenities, leading to the well-known issue of ability bias. In fact, a recent review by Mas (2025) concludes omitted productivity measures are "a primary source of bias in compensating differential estimates."

[1] This note handles the case of a single amenity, but the model readily extends to pricing multiple amenities. Researchers looking to price multiple amenities together should see the `aivreg` package or the multiple-amenity extension specified by Bell (2020).

[2] The more general case is handled by Bell (2020), where workers also get better or worse jobs for more idiosyncratic reasons, like search frictions, luck, or discrimination. While these are important extensions to make this model more in line with modern thought on wage determination, excluding them here changes nothing about the estimation strategy.

The solution involves two regressions:

$$h_i = \pi_w w_i + \pi_z z_i + \nu_i \tag{1}$$

$$w_i = \alpha + \beta z_i + \hat{h}_i + \varepsilon_i \tag{2}$$

where $\hat{h}_i$ in Equation 2 is the predicted value from Equation 1. Under the single-index assumption, $\hat{h}_i$ is the empirical analogue of η, and so equation 2 thus yields an estimate of β that is as good as if we could control for true ability η.

To illustrate why this works, take the simple (though restrictive) case where h is just a noisy version of ability; in other words, suppose h is η plus an iid disturbance.[3] Because noise in the outcome does not change regression coefficients, Equation 1 would have the same coefficients as if η were the outcome; the predicted values from this regression, then, re-construct η, which is the control that we wish to include in Equation 2.[4]

An alternative estimation strategy is to estimate β as $-\pi_z/\pi_w$ from Equation 1, much like taking the marginal rate of substitution from a production function. Both estimators are mathematically equivalent, and both estimators incorporate the insight that productivity should be in the residual of Equation 1.

Because the estimator can be written as a ratio of coefficients, a similar weak inference problem to IV can apply. The freely available `aivreg` Stata package provides Anderson-Rubin inference options that are robust to finite sample bias. Additional packages for non-parametric estimation are also available; the proof of non-parametric identification is contained in Bell (2020). At a basic level, non-parametric estimation of compensating differentials can be achieved through non-parametric estimation of Equations 1 and 2.

The core remaining limitation of the method, like earlier hedonic methods using observational data, is that it requires the researcher to observe the relevant set of amenities that contribute to workers' utility; we still cannot price amenities that we cannot see. However, the benefit in relation to prior methods is that the researcher no longer needs to observe and control for *all* of the ability channel.

[3] We entertain a measurement error assumption here solely to build intuition. This paragraph aside, the general model to have in mind is that many factors lead to better jobs, and h is just one that we observe.

[4] In this example where h is a noisy measure of η, controlling for h instead of $\hat{h}$ in Equation 2 leads the estimate of β to be biased — the typical "ability bias" problem. Importantly, as h tends toward a noisy measure of η, that bias grows. In contrast, using my method, increasing noise in h leads to increased *variance* in estimating Equation 1, and therefore a less precise — though *unbiased* — estimate of β.

III. Assumptions

The anti-instrument method for pricing job amenities is characterized by two key assumptions,[5] which distinguish it from other known approaches to this research question:

- *Single index:* If two workers share the same observed data on wages and amenities, then they must have the same latent η.
- *Conditional independence:* If we could observe and control for latent η, observed h would be redundant; i.e., h would have a 0 coefficient when regressing wages on amenities with η as a control.

An economic model of compensating differentials that rationalizes why the single index assumption is likely to hold was provided by Bell (2020). Intuitively, omitted factors that lead to better earnings–"ability," education, various skills, luck, search frictions, etc.–are also generally thought to yield better amenities.[6] Therefore, subsequent research applying the AIV method to job amenity pricing generally focuses on evaluating the conditional independence assumption as it relates to the specific choice of anti-instrument and amenities. Like the IV exogeneity assumption, the AIV conditional independence assumption can also be tested in an over-identified GMM framework (e.g., as implemented by the `aivreg` command).

Furthermore, these assumptions may be generalized to other research questions than job amenities. In that case, researchers must also establish the validity of the single index assumption for a new problem.[7]

IV. 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

[5] Bell (2020) states these assumptions more formally, and also discusses a third regularity condition pertaining to monotonicity.

[6] For this reason, although I have loosely referred to η as "ability," a more accurate term would be "job quality," to be inclusive of the myriad reasons why some workers get better jobs than others.

[7] In general, if a researcher wishes to use AIV to learn the effect of a variable (or vector) X on an outcome Y, both key assumptions must be evaluated. Conditional independence entails finding an anti-instrument that would be redundant with some latent variable η. This is almost always doable, because it is analogous to finding a single good control. The more challenging hurdle is to establish–with the help of economic reasoning and subject matter expertise–that any observation containing the same observed X and Y must also share the same value of that latent variable η. Not all questions can be well approximated by a single-index residual in this way. Research questions that are not well approximated by a single-index residual are not suitable to be estimated by AIV, but may still benefit from a popular and complimentary family of selection correction frameworks that involve alternative sets of assumptions to those presented here (e.g., see Oster (2019)).

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 educational attainment as an anti-instrument rather than as control (without any controls). The final two columns repeat this contrasting analysis, but when pricing a parsimonious set of amenities (one from each category) jointly in the same regression 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 `aivreg` Stata command.

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 reference 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 shows how applying the anti-IV estimator not only changes the sign of those original results, but also reconciles observational estimates with experimental estimates of workers' preferences.

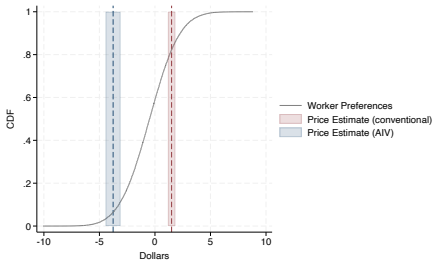
Mas and Pallais (2017) noted that their hedonic regressions produced counter-intuitive results: more pleasant work arrangements were associated with higher wages, contrary to the notion of compensating differentials. They write that "Table 1 shows the difficulty of estimating compensating differentials 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 empirical model delivers unexpected signs, but this approach is akin to regressing prices on quantities while controlling for all available supply shifters: the results of the regression may at face value seem to reject downward-sloping demand, but only because the model is mis-specified.

Table 1 presents the comparison. Column (1) reproduces the results presented in Mas and Pallais (2017, Table 1, Column 1), which attempts to control away ability using observed characteristics such as education, gender, and marital status. 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) applies the anti-IV estimator, using education as a noisy proxy for ability. In five of six cases, the signs flip: flexibility and work from home appear as costly amenities that workers are willing to sacrifice wages to obtain, while irregular schedules carry a wage premium. Pooling amenities within each category changes magnitudes but not signs. Throughout, the conventional approach and the anti-IV approach systematically deliver opposite conclusions.

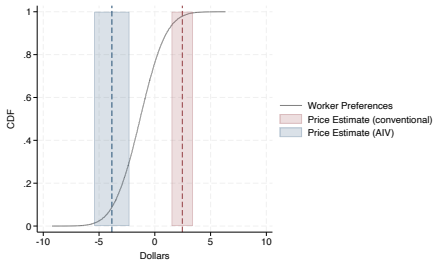
Figure 1 relates these last two columns of price estimates directly to the preference data elicited by Mas and Pallais (2017). Panel A shows that most workers positively value flexible schedules, yet the conventional model suggests wages

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

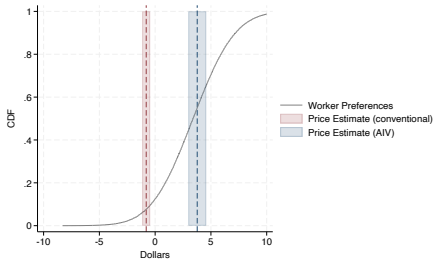
Figure 1. : Preferences and Prices for Amenities



(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 education as anti-instrument when relating hourly earnings to the given amenity, analogous to Column 4 of Table 1.

are higher for those with flexible jobs. In contrast, the anti-IV estimates align with preferences: workers give up pay to obtain flexibility, and the rarity of flexible jobs reflects their high cost to employers. Panel B reports a similar pattern for working from home. Workers state large willingness to pay for remote work, and the anti-IV estimate confirms it as an expensive amenity–explaining why it was scarce 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 anti-IV estimates instead show that workers require a pay premium to take irregular schedules, 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 anti-IV 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 ability bias undermines conventional hedonic regressions, and that experimental methods can overcome this ability bias. However, ability bias should not be regarded as a limitation inherent to cross-sectional observational data, as models endorsed by prior literature would imply. With the anti-IV estimator, the very same cross-sectional observational data by and large replicate and validate the insights of Mas and Pallais's experiment. 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 insights in non-experimental settings substantially expand their applicability.

V. Conclusion: A Practical Tool for a Common Problem

A growing number of applied papers have used this estimator 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). The estimator presented here offers an improvement over existing methods for estimating compensating differentials because it assumes unobserved productivity causes both wages and amenities, as opposed to only wages. It is simple enough for applied researchers, yet rigorous enough to address a core identification problem. By incorporating the reality that productivity is in the residual of wage regressions, and using the methods outlined in this note to control for it, researchers can estimate amenity prices in cross-sectional data without relying on natural experiments or strong structural assumptions.

This short note provides the essentials, while other work presents formal results, model extensions, and empirical applications. For implementation, the `aivreg` package offers an immediate way to get started. As with productivity estimation, treating ability as a residual—rather than a control variable—will continue to open new avenues for credible empirical work.

REFERENCES

Bell, Alex, "Job Amenities & Earnings Inequality," *MPRA Paper*, May 2020. Number: 118516 Publisher: University Library of Munich, Germany.

_ , 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. Publisher: INFORMS.

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

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.

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

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

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, April 2019, *37* (2), 187–204.

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.

Solow, Robert M., "A Contribution to the Theory of Economic Growth," *The Quarterly Journal of Economics*, February 1956, *70* (1), 65.