Dissertation Awards

2016

Ranking Firms Using Revealed Preference and Other Essays about Labor Markets

Isaac Sorkin
The University of Michigan

Ranking Firms Using Revealed Preference and Other Essays about Labor Markets

Isaac Sorkin

This dissertation contains essays on three questions about the labor market. In light of the fact that firms account for a substantial share of earnings inequality, Chapter 1 asks why some firms pay their employees so much and some so little. Although the standard explanation is that search frictions support an equilibrium with rents, this paper finds that compensating differentials are at least as important. To reach this finding, the paper develops a structural search model and estimates it on U.S. administrative data with 1.5 million firms and 100 million workers. The model analyzes the revealed preference information contained in how workers move between firms. Compensating differentials are revealed when workers systematically move to lower-paying firms, while rents are revealed when workers systematically move to higher-paying firms. With the number of parameters proportional to the number of firms (1.5 million), standard estimation approaches are infeasible. The paper develops an estimation approach that is feasible for data on this scale. The approach uses tools from numerical linear algebra to measure central tendency of worker flows, which is closely related to the ranking of firms revealed by workers' choices.

Chapter 2 asks why men and women work at different firms. It builds on Chapter 1 to consider two explanations: 1) men and women might search from different offer distributions, and 2) men and women might have different rankings of firms. I find that the first explanation—women search from a lower-paying offer distribution than men—is the dominant one.

Chapter 3 asks what are the long-run effects of the minimum wage? An empirical consensus suggests that there are small employment effects of minimum wage increases. This chapter argues that these are short-run elasticities. Long-run elasticities, which may differ from short-run elasticities, are more policy relevant. To shed light on these questions, this chapter develops a dynamic industry equilibrium model of labor demand. The model makes two points. First, long-run regressions have been misinterpreted because even if the short- and long-run employment elasticities differ, standard methods would not detect a difference using U.S. variation. Second, the model offers a reconciliation of the small estimated short-run employment effects with the commonly found pass-through of minimum wage increases to product prices.

Chapter 1: Ranking Firms Using Revealed Preference

Firms play a central role in explaining worker earnings. Conditional on person fixed effects, firms account for over 20 percent of the variance of worker earnings (e.g., Abowd, Kramarz, and Margolis [1999] and Card, Heining, and Kline [2013]). But there is little work asking why firms play such a central role. There are two main explanations that differ in whether, from the worker perspective, high-paying firms are high-value firms: rents and compensating differentials. Rents is the leading explanation in the literature (e.g., Postel-Vinay and Robin [2002]). In the rents explanation, frictions prevent workers from bidding away the rents at high-paying firms. In this explanation, workers are lucky to end up at high-paying firms. As a result, high-paying firms are high-value firms. In contrast, in the compensating differentials explanation (e.g., Rosen [1986)]), firms differ both in how much they pay and in nonpay characteristics. In this explanation, higher pay compensates for variation in unpleasant nonpay characteristics. As a result, high-paying firms are not high-value firms. Thus, measuring the relative importance of rents and compensating differentials means figuring out whether high-paying firms are high-value firms.

To distinguish between high-paying and high-value firms, this chapter estimates the value of working at a firm without using information in pay. Specifically, I use information in quantities. To do so, I exploit the revealed preference idea embedded in search models that workers move to firms with higher value. Using U.S. matched employer-employee data with 1.5 million firms, I map the 1.5 million by 1.5 million matrix of worker flows across firms into the value of working at each firm.

By comparing the firm-level values to firm-level pay, I find that both rents and compensating differentials explanations are operative, but compensating differentials are more important. This finding has four (closely related) implications. First it shows that in many cases the conventional interpretation of high-paying firms as "good" firms is not warranted. Second, it contrasts with the conventional wisdom that compensating differentials are unimportant in explaining the variance of worker earnings. Third, it resolves a puzzle that benchmark search models are unable to reproduce the extent of earnings dispersion while also yielding plausible values of nonemployment. Fourth, the distribution of compensating differentials across firms means that the variance of earnings is larger than it would be if all jobs were equally pleasant.

In the first part of the chapter, I develop a simple model of the labor market that contains both the rents and the compensating differentials explanations and develop tools to estimate it using only quantity information. The model is a benchmark partial equilibrium utility-posting model in the spirit of Burdett and Mortensen (1998). The nonstandard

ingredients in the model are that firms post a utility offer that consists of both earnings and a nonpecuniary bundle and that workers receive idiosyncratic utility draws (taste shocks) each period. On the one hand, the rents explanation is contained in the model because there is the possibility of equilibrium dispersion: different firms offer different levels of utility. On the other hand, the compensating differentials explanation is contained in the model because high earnings might be offset by a low nonpecuniary bundle. The role of the idiosyncratic utility draws is to add preference heterogeneity and allow workers to make different decisions when faced with the same choice. Specifically, these draws explain why there would be flows both from firm A to firm B and from firm B to firm A.

I estimate the model in two steps. First, I isolate the transitions that reveal preferences by using information about what the worker's coworkers were doing at the time of the separation. In the spirit of the displaced worker literature (Jacobson, Lalonde, and Sullivan 1993), if an unusually high share of coworkers were also separating, then a firm-level shock caused the workers to leave, and there is a high probability that any particular separation was exogenous. In contrast, if turnover levels look normal, then this separation was likely idiosyncratic to the worker.

Second, I measure the central tendency of worker flows. I record worker flows in a 1.5 million by 1.5 million matrix, where one cell is the number of workers who go from firm A to firm B. The model implies a set of linear restrictions on the entries in this matrix and a flow-relevant firm-level value. The flow-relevant firm-level value captures the central tendency of worker flows and is a known function of structural parameters.

Computing this central tendency of worker flows—and showing that it exists and has a meaningful economic interpretation—is the main technical contribution of this chapter. The central tendency of worker flows is captured by the top eigenvector of a normalized matrix of worker flows. Showing when this eigenvector exists and is unique requires a new analytical result. Computing the eigenvector relies on techniques from numerical linear algebra that are scalable to massive data sets such as a 1.5 million by 1.5 million matrix.

In addition to the value of a firm, two factors affect the central tendency of mobility. First, a large firm will naturally have more workers moving away from it than a small firm. I account for this because I observe firm size. Second, a firm that makes a lot of offers will naturally have more workers moving toward it. I account for this because I estimate the offer distribution using information in nonemployment-to-employer flows. By jointly estimating the offer distribution and the value of nonemployment, I allow nonemployed workers to reject offers.

In the second part of the chapter, I estimate the earnings that firms post and show how to combine these with the estimates of firm values to decompose the variation in firm-level earnings into compensating differentials and rents. The model implies that—as in Abowd and Schmutte (2014)—I can estimate the earnings that firms post using a selection-corrected version of the statistical decomposition pioneered by Abowd, Kramarz, and Margolis (1999) (and also used by Card, Heining, and Kline [2013]), which controls for person effects.

Because I use revealed preference at the firm level, I can identify compensating differentials. At the individual level, about a third of employer-to-employer moves result in earnings cuts. As the literature recognizes, this finding does not identify compensating differentials because the pay cut might reflect idiosyncratic factors that are not priced in the labor market. At the firm level, however, systematic patterns of workers moving from higher-paying to lower-paying firms indicates factors that are valued by all workers and are priced in the labor market.

Formally, I prove an identification result to show how to measure the relative roles of compensating differentials and rents. Combining utilities and earnings gives a lower bound on the variance of nonpecuniary characteristics, which is the extent of compensating differentials. The complement is the role of rents in explaining the variance of earnings. The identification result is consistent with an argument in the search literature that compensating differentials are hard to find because desirable nonearnings factors might be positively correlated with earnings (Hwang, Mortensen, and Reed 1998; Lang and Majumdar 2004). Specifically, I distinguish between Rosen and Mortensen amenities. Rosen amenities generate compensating differentials. In contrast, Mortensen amenities are positively correlated with earnings and generate equilibrium dispersion; for example, some high-paying firms might also offer great benefits. The result shows that in my framework the variance of Mortensen amenities is not identified.

I estimate the model on the U.S. Census Bureau's Longitudinal Employer Household Dynamics data set. Compensating differentials explain twice as much of the variance of firm-level earnings as rents. The estimated ranking of sectors is intuitively plausible, as is the implied distribution of nonpay characteristics. For example, education has good nonpay characteristics, while many blue-collar sectors have bad nonpay characteristics. The central finding that compensating differentials are at least as important as rents holds within subgroups defined by age and gender.

The finding that compensating differentials are important contrasts with conventional wisdom.[1] Research that looks at specific amenities rarely finds robust evidence that amenities are priced in the labor market. By using revealed preference at the firm-level, however, I overcome a few challenges in this literature. First, I distinguish between *Rosen* and *Mortensen* amenities to allow for some amenities to be unpriced and contribute to dispersion in firm-level value. Second, revealed preference takes into account the entire

bundle of amenities. As a result, it does not take a stand on what workers value or how firms use amenities to offset other amenities. Third, at the firm level, revealed preference uncovers amenities that are priced. In contrast, at the individual level, revealed preference can uncover idiosyncratic factors that are unlikely to be priced.

The finding of a large role for compensating differentials helps resolve the puzzle emphasized by Hornstein, Krusell, and Violante (2011) that benchmark search models cannot generate the extent of observed residual earnings inequality. I find that workers act as if a large share of the variance of firm-level earnings does not reflect variation in value. I also find an empirically reasonable value for nonemployment, and thus pass a crucial test they propose.

Finally, if the estimated nonpay characteristics were removed and earnings changed to compensate workers, then earnings inequality would decline. The effect depends on the correlation between earnings potential and nonpay characteristics. In this counterfactual, earnings inequality as measured by the variance of earnings declines. Many low-paying jobs come with good amenities, so this reduction comes mainly from the lower tail of the income distribution shifting up.

Chapter 2: Why Do Men and Women Work in Different Firms?

Men are more likely than women to work in both high-wage firms and high-wage industries. This sorting component accounts for 9 log points of the 36 log point (about 25 percent) gender earnings gap in the United States. A core interpretive issue is whether this sorting component reflects discrimination or differences in preferences.[2] The discrimination explanation is that women would like to work at the same firms or in the same industries as men, but are prevented from doing so. The preference explanation is that there are nonpecuniary characteristics that differ across the low- and high-paying firms, and women value these nonpecuniary characteristics more than men. For example, women might highly value flexibility while men do not.

In this chapter, I build on Chapter 1 to shed light on these two explanations. I begin by establishing the key fact that this chapter is interested in understanding: men are at higher-paying firms than women in the United States, and more generally, that men and women are sorted in the labor market. Building on Card, Cardoso, and Kline (2016), who study Portugal, I estimate the Abowd, Kramarz, and Margolis (1999) decomposition separately for men and women using data from the United States. This decomposition amounts to using the complete set of wage changes of workers who switch firms to estimate a firm effect in earnings. Estimating the decomposition separately for men and women allows me to construct gender-specific earnings premia at each firm. I reproduce the finding of Card, Cardoso, and Kline (2016) that the sorting of men and women across firms is quantitatively important; indeed, I find that the sorting component is more important in the United States than in their Portugese context.

I then explain why men and women are sorted. Following Chapter 1, I record a partial equilibrium utility-posting search model in the spirit of Burdett and Mortensen (1998), where the only nonstandard ingredient is that firms post utility offers that combine a wage and a nonpecuniary bundle, and there is a (transitory) idiosyncratic utility draw in each match, which explains why people might make different choices. The model estimates revealed values of employers by taking into account the network structure of accepted offers and rejected offers.

The model embeds versions of both the preference and discrimination explanations for why men and women are sorted. The preference explanation in the model is that men and women rank firms differently; therefore, given the same set of opportunities, they would end up in different firms. The discrimination explanation in the model is that men and women receive a different set of offers, and so given the same preferences end up in different firms.

To separate the preference and discrimination explanations, I estimate the model separately by men and women. The implicit assumption in this exercise, as with standard earnings decomposition exercises, is that men and women operate in separate labor markets. This allows me to estimate separate offer distributions (opportunities), values (preferences), and earnings for men and women. In particular, I am able to estimate separate offer probabilities, values, and earnings for men and women firm by firm. While the model relies on numerous strong assumptions, estimation is completely nonparametric along the dimensions that this chapter is interested in. Namely, I impose no assumptions on the shape of the offer distribution, no assumptions on the distribution of employer values, and no assumptions on the relationship between men's and women's preferences or offer distributions.

My principal results are as follows. First, I find that men and women are systematically sorted in the labor market. About 60 percent of men's coworkers are men, while only about 40 percent of women's coworkers are men. This finding is quantitatively consistent with the U.S. results in Hellerstein, Neumark, and Mcinerney (2008, p. 183). Second, I find that this sorting explains over 25 percent of the earnings gap (given the nature of the data, some of this might include differences in hours). This finding is robust to using men's or women's earnings to compute how well-paying firms are. This extends the finding of Card, Cardoso, and Kline (2016) (for Portugal) that men are systematically employed in higher-paying firms than women.

I then use the model to interpret the reasons and consequences of sorting. The model points to differences in the offer distribution rather than differences in preferences to explain why men and women are sorted. Men's and women's

preferences over firms are estimated to be highly correlated. For example, the overall estimated values have a correlation of 0.89. When I aggregate to the 4-digit industry level (which explains two-thirds of the sorting component of the gender wage gap), the correlation is 0.98. But I estimate that women search from a lower-paying distribution. This result should be interpreted cautiously since the offer distribution is a reduced-form object that may itself contain revealed preference information. (Put differently, the fact that women are less likely to receive offers from high-paying firms may reflect the fact that they do not want to work at those firms, rather than the fact that they do not receive offers from these firms.) The model also allows me to compare the distribution of employer values at the firms that employ men relative to the firms that employ women. Taking the model at face value, the estimates indicate that men and women are at approximately equal-valued firms regardless of whether it is men's or women's values that are used to value the firms.

Chapter 3: Are There Long-Run Effects of the Minimum Wage?

Inflation and rising real wages make most minimum wage increases temporary. As such, the empirical minimum wage literature has made substantial progress estimating the short-run employment effects of minimum wage increases. This effect appears to be small. Despite apparent consensus, the profession remains divided about the employment effects of minimum wage increases.

A reasonable reading of this divide is that there are some questions about the effects of minimum wage increases for which the empirical consensus provides the answer. For other questions, however, economists extrapolate differently depending on whether they think that the relevant short- and long-run employment elasticities differ. To the question, "What is the employment effect of a temporary nominal minimum wage increase likely to be?", the empirical consensus suggests that there are unlikely to be significant employment effects, because similar increases have not resulted in significant employment effects. To the question, "What is the employment effect—after a few years—of a permanent minimum wage increase?", the empirical consensus suggests an answer only if the short- and long-run elasticities of minimum wage increases are the same. In the United States, this latter question is of immediate policy relevance: President Obama's 2013 State of the Union address contained a proposal to index the federal minimum wage to inflation, which would be a more permanent increase.

To contribute to this important debate, this chapter studies the empirical implications of a model that distinguishes between the short- and long-run employment elasticities. The model is based on the putty-clay nature of capital. It was first informally discussed in the minimum wage context by Card and Krueger (1995, pp. 366–368), and I build on the Gourio (2011) version.

In the model, when firms pay the entry cost of building a machine, they can freely substitute between capital and labor. Once capital is installed, a firm cannot change its labor demand.

The key features of the model are that the labor demand choice of an entering firm is a forward-looking, dynamic decision that depends on the (expected) stochastic process for minimum wages. And because only some firms adjust each period, the industry-level labor demand response to a minimum wage increase is slow, and also depends on the stochastic process for minimum wages.

The model has two main empirical implications. The first empirical implication is that the reduced-form long-run effects estimated in the literature are essentially uninformative about the true long-run elasticity. I simulate employment data from the model to replicate the data set used in Dube, Lester, and Reich (2010), who find very small short-run employment effects and, using a common reduced-form long-run regression, find no distinction between the short- and long-run employment effects of minimum wages in the United States. They interpret these results as evidence against the view that short- and long-run elasticities differ. On the simulated data, however, the reduced-form regression recovers a long-run employment effect that is barely different than the short-run employment effect. As such, the chapter suggests that it would be a mistake to infer from existing empirical work that proposals to index minimum wages to inflation would have minimal effects on employment. Taking the model at face value shows how misleading such an inference might be: the results in the chapter show that a contemporaneous elasticity of -0.002 in response to a temporary increase is consistent with an elasticity after six years of -0.252 for a permanent increase.

The second empirical implication is that the putty-clay model is consistent with the pass-through of minimum wage increases to product prices commonly found in the literature, even though minimum wage increases are relatively temporary. Card and Krueger (1994, p. 792) emphasize that their finding of product price rises in response to minimum wage increases are inconsistent "with models in which employers face supply constraints (e.g., monopsony or equilibrium search models)." Despite this, the minimum wage literature has focused on models of search frictions to rationalize the small employment effects without focusing on the price results.

Notes

1. For example, Hornstein, Krusell, and Violante (2011, p. 2883) write that compensating differentials "does not show too much promise" in explaining earnings dispersion.
2. See, for example, the Council of Economic Advisers (2015, p. 3). https://www.whitehouse.gov/sites/default/files/docs/equal_pay_issue_brief_final.pdf (accessed October 6, 2016).

References

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

Abowd, John M., and Ian M. Schmutte. 2014. "Endogenous Mobility." Unpublished working paper.

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

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics* 131(2): 633-686.

Card, David, Joerg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128(3): 967–1015.

Card, David, and Alan B. Krueger. 1994. "Minimum Wages and Employment: A Case Study of the Fast Food Industry in New Jersey and Pennsylvania." *American Economic Review* 84(4): 772–784.

———. 1995. *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton, NJ: Princeton University Press.

Dube, Arindrajit, T. William Lester, and Michael Reich. 2010. "Minimum Wage Effects across State Borders: Estimates Using Contiguous Counties." *Review of Economics and Statistics* 92(4): 945–964.

Gourio, Francois. 2011. "Putty-Clay Technology and Stock Market Volatility." *Journal of Monetary Economics* 58(2): 117–131.

Hellerstein, Judith, David Neumark, and Melissa Mcinerney. 2008. "Changes in Workplace Segregation in the United States between 1990 and 2000: Evidence from Matched Employer-Employee Data." In *The Analysis of Firms and Employees: Quantitative and Qualitative Approaches*, Stefan Bender, Julia Lane, Kathryn Shaw, Fredrik Andersson, and Till von Wachter, eds. Chicago: University of Chicago Press, pp. 163–195.

Hornstein, Andreas, Per Krusell, and Giovanni L. Violante. 2011. "Frictional Wage Dispersion in Search Models: A Quantitative Assessment." *American Economic Review* 101(7): 2873–2898.

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

Jacobson, Louis, Robert LaLonde, and Daniel Sullivan. 1993. "Earnings Losses of Displaced Workers." *American Economic Review* 83(4): 685–709.

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): 1111–1128.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70(6): 2295–2350.

Rosen, Sherwin. 1986. "The Theory of Equalizing Differences." In *Handbook of Labor Economics. Vol. 1*, Orley Ashenfelter and Richard Layard, eds. Amsterdam: Elsevier, pp. 641–692.