

2014

Essays on Labor Market Inequality: Dissertation Summary

Conrad Miller
MIT

Essays on Labor Market Inequality

Conrad Miller

This thesis consists of three chapters on aspects of labor market inequality. Chapter 1 examines the dynamic effects of federal affirmative action regulation, exploiting variation in the timing of regulation and deregulation across work establishments. Chapter 2 studies the spatial mismatch hypothesis, which proposes that job suburbanization isolates blacks from work opportunities and depresses black employment. Chapter 3, coauthored with Isaiah Andrews, analyzes the effect of heterogeneity on the widely used analyses of Baily (1978) and Chetty (2006) for optimal social insurance.

Chapter 1

The Persistent Effect of Temporary Affirmative Action

The first chapter analyzes the dynamic effects of temporary federal affirmative action regulation. Affirmative action policies—those designed to increase diversity among employees, students, politicians, or businesses by advantaging candidates from underrepresented social groups—are practiced throughout the world (Fryer and Loury 2013). They are universally controversial. Even among their advocates, they are often introduced or supported as only temporary remedies for existing social inequities (Sowell 2004). The hope is that a temporary affirmative action program that enhances diversity and reduces inequality between groups can persistently alter those outcomes.

Whether a temporary policy will indeed have persistent effects remains an open question. Economic theory provides ambiguous predictions. The theoretical literature primarily focuses on the potential for affirmative action to reduce inequality by incentivizing human capital accumulation for disadvantaged groups (e.g., Lundberg and Startz 1983). If employers perceive that some group of workers is less productive or have more difficulty screening workers from that group, then the return to human capital investment for members may be inefficiently dampened. In this setting, an affirmative action regulation can correct those incentives, and even a temporary program can permanently reduce inequality by eliminating negative stereotypes. However, Coate and Loury (1993) demonstrate this need not be the case; indeed, affirmative action can reduce the return to investment for that group even further. In this case, an affirmative action policy must be maintained permanently for any protected group gains to persist. In general, the consequences of a temporary affirmative action policy may depend on the setting.

In this chapter I study the dynamic effects of Executive Order 11246, the primary affirmative action regulation for

employment in the United States. The regulation applies to firms that have sizable contracts or subcontracts with the federal government. The Department of Labor estimates that such firms employ about a quarter of the U.S. workforce (Office of Federal Contract Compliance Programs 2013). Regulated firms are mandated to make a “good faith” effort to employ minorities at rates (at least) proportional to their shares of the local and qualified workforce. I study the regulation’s effect on the employment of black workers, one of the regulation’s original targets, the largest minority group over my period of study, and a group that is often the focus of affirmative action research (Holzer and Neumark 2000a). My work builds on the influential analysis of Leonard (1984), and the recent, careful, and closely related studies of the impacts of Executive Order 11246 on the employment and occupational advancement of women and minorities by Kurtulus (2011, 2012).

To estimate the dynamic effects of federal affirmative action regulation, I use an event study research design, exploiting variation in the timing of regulation and deregulation across work establishments. In particular, I utilize changes in employers’ status as a federal contractor using administrative data from 1978 to 2004. For many types of goods and services, the set of companies the government buys from at any given time is constantly changing. Turnover in these contractor relationships provides plausibly exogenous variation in which and when employers are subject to affirmative action regulation.

I find that affirmative action sharply increases an establishment’s share of employees that are black, with the share continuing to increase over time. Five years after an establishment is first subject to the regulation, its share of employees that is black increased by an average of 0.8 percentage points. To put this magnitude in perspective, note that a 0.8–1.3 percentage point increase in the share of the U.S. workforce that is black would eliminate the black-white jobless gap over this period. This effect is proportionally larger for middle- and high-skill occupations.

Strikingly, I find that the share of employees that is black continues to grow even after an employer is deregulated. This persistence is evident more than a decade following deregulation. By contrast, gaining and losing contractor status have symmetric associations with other employer characteristics. Establishment size increases when an establishment becomes a contractor, and decreases when it loses its contractor status. Moreover, following deregulation, an establishment’s likelihood of acquiring a new contract—and hence, becoming regulated again—quickly reverts to near the baseline rate.

This persistence is difficult to reconcile with existing economic models of affirmative action, which focus on the aforementioned human capital channel (Fang and Moro 2011). In particular, because the policy variation exploited here varies across individual employers, it should have minimal effects on the human capital investment incentives work-

ers face in the broader labor market. Rather, any response is likely driven by changes at the employer level.

Given that employers continue to increase the share of their workforce that is black even when they are deregulated, a revealed preference argument would imply that it is profitable for them to do so. Consistent with this, I argue that the persistence found here is in part due to employers investing in what I term *screening capital*—investments that improve an employer’s ability to screen potential workers. Examples may include employing and training personnel specialists and departments, developing job tests, harnessing referral networks, developing relationships with and utilizing intermediaries such as employment agencies and schools, and even learning by doing or experimentation (Arrow 1962; Fryer and Jackson 2008).¹

Building on the seminal Phelps (1972) model of statistical discrimination, I show how the persistence found here may be driven by affirmative action inducing employers to make (partially) irreversible investments to improve screening. In existing models, an employer can only comply with affirmative action by reducing their hiring standard for the protected group. I introduce a novel response margin, allowing employer investments in screening capital. I show that, under conditions often assumed in the statistical discrimination literature, screening investments will reduce between-group disparities in hiring rates; moreover, affirmative action will increase the return to such investments. If these investments are at least partially irreversible, temporary affirmative action regulation can generate persistent changes in screening capital, and hence produce a durable increase in the minority share of hires.

I then present evidence supporting the model’s predictions. First, the model predicts that regulation will increase the return to investments in screening. Using cross-sectional survey data, I show that regulated employers use more screening methods—including personnel specialists, job tests, credential checks, and intermediaries—than otherwise similar unregulated employers. These results largely echo those in Holzer and Neumark (2000b).

Second, the model predicts that screening investments will reduce between-group differences in hiring rates. To test this, I exploit another source of variation in screening investment: employer size. It is well documented that larger employers use more resources in screening and use a wider variety of methods (Marsden 1994). Using administrative panel data, I show that employers’ share of employees that is black is increasing in employer size. While previous work documents a positive cross-sectional correlation between employer size and share of employees that are black (Carrington, McCue, and Pierce 2000; Holzer 1998), it is not clear what drives this relationship. For example, the authors of these studies posit that the relationship may be driven by workplace discrimination law, which does not cover establishments with fewer than 15 employees, or the concentra-

tion of larger establishments in urban locations. To rule out these alternative explanations, I show that this relationship holds within-establishment and within-job (where jobs are defined as establishment by occupation cells) for a large sample of establishments that are all subject to workplace discrimination law.²

Given that employers continue to increase the share of their workforce that is black even when they are no longer regulated, a revealed preference argument would imply that it is profitable for them to do so. This suggests that affirmative action leads firms to take actions that increase the profitability of employing black workers. I provide evidence that these actions include screening investments. These investments might take several forms, and determining what types of capital are most significant for both compliance and persistence requires further study. Alternatively, affirmative action may prompt employers to change their personnel practices in a way that is prohibitively costly to reverse. To distinguish between these classes of explanations, one could measure the dynamic effects of affirmative action on productivity and profitability. More generally, the persistence documented in this chapter suggests that minority workers face job search frictions that can be at least partially surmounted by temporary intervention. Understanding the sources of these frictions and specific mechanisms that can mitigate these impediments remains an important area for future research.

Chapter 2

When Work Moves: Job Suburbanization and Black Employment

The second chapter introduces novel evidence on the spatial mismatch hypothesis. Since 1960, the unemployment rate in the United States for blacks has been roughly double the unemployment rate for whites (Fairlie and Sundstrom 1999). Large disparities remain after conditioning on measures of labor market skill, including education and Armed Forces Qualification Test scores (Ritter and Taylor 2011). The spatial mismatch hypothesis argues that this gap is in part attributable to the geographic distribution of employers and households. Black households tend to live relatively far from work opportunities, reducing their access to employment. In particular, while firms and white households relocated from the central city to the suburban ring at an accelerated rate following World War II—from 50 percent of jobs and white residents in 1950 to 30 percent in 1990 (Baum-Snow 2007)—black households faced initially strong barriers to suburban residence, including housing discrimination and liquidity constraints. As a result, they remained concentrated in the central city. Introduced by John Kain (1968), the spatial mismatch hypothesis and related ideas were further popularized by sociologist William Julius

Wilson in *The Truly Disadvantaged* (1987) and *When Work Disappears* (1996).

An extensive literature sets out to test the spatial mismatch hypothesis, typically by relating labor market outcomes to measures of job accessibility in a cross-section.³ Most recent work in this literature finds some support for spatial mismatch, though results tend to be sensitive to how job accessibility is measured.⁴ More importantly, results from this literature are made difficult to interpret by the endogeneity of household and firm location. Across and within metropolitan areas, residents who are (unobservably) less productive may sort into neighborhoods that are farther from work opportunities, where rents are typically lower.⁵ Similarly, firms may choose to locate in neighborhoods with residents who are (unobservably) more productive. In this chapter I test the spatial mismatch hypothesis over several decades while accounting for these endogenous location decisions.

In my analysis, I account for the endogeneity of household and firm locations in two steps. First, to account for household sorting I construct synthetic cohorts and estimate models in differences, exploiting variation in job suburbanization across metropolitan areas. In particular, I estimate models relating changes in cohort employment rates to changes in the spatial distribution of work. Surprisingly, panel methods have rarely been applied in this setting.⁶ This approach has two advantages: 1) by conditioning on baseline employment, I absorb variation in time-invariant unobservable characteristics (as well as the effects of any initial spatial mismatch); and 2) it allows for households to sort across neighborhoods. Of course, with synthetic cohorts one concern is that endogenous migration may introduce compositional changes so that the relationship between job suburbanization and employment growth may in part reflect the changing composition of synthetic cohorts rather than within-person changes.⁷ Fortunately, I am able to test directly for endogenous migration responses and find that any compositional changes are negligible.

Using census data from 1970 and 2000, I find that job suburbanization is associated with significant relative declines in black employment. For every 10 percent decline in the fraction of MSA jobs located in the central city over this period, black employment rates declined by 1.4 to 2.1 percent relative to white employment. Relative earnings declined by 1.1 to 2.3 percent. These relationships are not artifacts of selective migration, residential suburbanization, or changes in industry or occupation composition. Consistent with the focus of the literature, this is a low-skill and middle-skill phenomenon; there is no relationship between job suburbanization and relative employment for blacks among college graduates. Notably, the estimates are driven almost entirely by job suburbanization that occurred during the 1970s. From 1970 to 1980, for every 10 percent decline in the fraction of MSA jobs located in the central city over this period, black relative employment rates declined by

about 2.6 percent. From 1980 to 2000, such suburbanization is associated with only a 0.5–0.6 percent decline in black relative employment rates. Given the magnitude of the estimates over the full period, this suggests that suburbanization occurring during the 1970s had a persistent effect.

Second, to account for the potential endogeneity of job suburbanization—in particular, changes in the spatial distribution of work driven by unobserved productivity shocks that are unevenly distributed across black and white labor—I exploit variation in job suburbanization that is plausibly exogenous to such productivity shocks. Specifically, I use variation in interstate highway construction across MSAs, as identified in Baum-Snow (2007). These highways were planned in the 1940s and 1950s, and hence their assignment across MSAs should be orthogonal to 1970s' residual labor supply shocks. I show that they appear to be unrelated to demand shocks particular to black workers over this period. In addition, as in Baum-Snow (2007), I use a 1947 federal plan for the interstate highway system that was explicitly designed to link far away places rather than facilitate local commuting or economic development to instrument for actual highways constructed. I test whether the causal impacts of highways on the labor market are consistent with spatial mismatch; in particular, if highways increase job suburbanization and reduce relative employment for blacks. To the best of my knowledge, this is the first work to exploit variation in highway construction to test the spatial mismatch hypothesis.

Using exogenous variation in highway construction, I find that highways predict job suburbanization and declines in black relative employment in a manner consistent with spatial mismatch. While highways cause suburbanization that continues with each decade, the decrease in relative employment for blacks emerges during the 1970s and stagnates thereafter. This suggests that the estimated relationship between job suburbanization and black employment is causal and not driven by unobserved shocks to worker productivity. In addition, I find that highways cause residential suburbanization of white households but not black households over my period of study, consistent with the premise that black households faced significant additional barriers to suburban residence. Altogether, the results suggest that job suburbanization was an important determinant of black labor market outcomes over the 1970s, and that its initial impact persisted.

One additional concern with my approach is that if job suburbanization is driven by the relatively high exit of firms in the central city—which may disproportionately affect black workers, given that they tend to be concentrated in the central city—the estimated negative relationship between job suburbanization and black employment may simply reflect the effects of job displacement rather than the spatial distribution of work per se. Using unique establishment data from the Equal Employment Opportunity Commission, I show that this is unlikely to be the case.

Chapter 3

Optimal Social Insurance with Heterogeneity

The third chapter, written with Isaiah Andrews, investigates the implications of population heterogeneity for the sufficient statistic approach to welfare analysis developed in Baily (1978) and generalized in Chetty (2006). In a stylized model of unemployment, Baily obtains a simple formula for the optimal unemployment insurance (UI) benefit as a function of three parameters: 1) the elasticity of unemployment durations with respect to benefits, 2) the drop in consumption associated with unemployment as a function of UI benefits, and 3) the coefficient of relative risk aversion.⁸ This framework has been applied extensively in both empirical and theoretical work on social insurance (e.g., Gruber 1997). One potential shortcoming of the Baily and Chetty results is that they are derived using models where agents are homogeneous along some important dimensions, while in practice heterogeneity seems likely to be empirically relevant. In the UI context, for example, there may be heterogeneity across agents in search costs; ability to smooth consumption (e.g., via borrowing, savings, or spousal labor supply); and local risk aversion. This heterogeneity can affect how individuals value UI, and the need to aggregate heterogeneous individual preferences may significantly complicate welfare analysis.

As noted by Chetty (2006), the Baily-Chetty formulas are robust to a limited degree of heterogeneity, provided one plugs in appropriate population averages.⁹ This result, however, requires the assumption that agents share a common coefficient of relative risk aversion. This homogeneity assumption is used to relate differences in average utility across states to differences in average consumption. By considering the joint distribution of risk aversion and consumption drops, we extend the Baily-Chetty framework to allow for arbitrary heterogeneity in risk preferences, and hence unrestricted heterogeneity across agents.

We show that several different approaches to calculating aggregate welfare for heterogeneous agents yield equivalent welfare metrics. We find that heterogeneity in risk aversion affects welfare analysis through the covariance between risk aversion and consumption drops in the cross-section of the unemployed. This reflects the fact that UI is more valuable if more risk-averse agents are subject to larger risks. We refer to this as the covariance effect.¹⁰ Our approach easily generalizes to accommodate a number of extensions, including UI systems with taxes and benefits that are proportional to wages. Further, we show that our results extend to a heterogeneous version of the rich dynamic model studied by Chetty (2006), allowing for a range of additional behaviors and constraints, including private insurance purchases and limits on borrowing.

To explore the potential importance of the covariance effect, we calibrate a stylized model of private consumption

smoothing decisions using data on observed household consumption drops associated with unemployment. The results suggest that the covariance effect may be large: for plausible population distributions of risk preferences, we find that accounting for the covariance effect can change the approximate consumption smoothing benefit of UI by more than 50 percent.

Our results show that the value of social insurance depends on the extent to which risk exposure and risk tolerance are aligned in the unemployed population: for a given distribution of consumption drops, the lower the covariance of consumption drops faced by workers with individual risk aversion, the lower the value of additional social insurance. In contexts where risk aversion, ex ante risk, and ability to self-insure are largely independent, we would generally expect this covariance to be negative because more risk-averse agents will take private actions to reduce their risk. Moreover, we would expect the magnitude of this effect to be larger when workers are better able to self-insure. To take an extreme example, even if most agents are quite risk averse and the average consumption drop associated with unemployment is large, the marginal value of social insurance may be zero if all consumption risk is borne by a risk-neutral subpopulation, as could occur in the presence of actuarially fair private UI. However, without knowing the joint distribution of risk preferences, ex ante risk, and ability to smooth consumption, the sign and magnitude of the covariance effect are a priori ambiguous and may depend on context; estimating this covariance is an important challenge for future research.

References

- Arrow, K. J. 1962. "The Economic Implications of Learning by Doing." *Review of Economic Studies* 29(3): 155–173.
- Baily, M. N. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10: 379–402.
- Baum-Snow, N. 2007. "Did Highways Cause Suburbanization?" *Quarterly Journal of Economics* 122(2): 775–805.
- Carrington, W. J., K. McCue, and B. Pierce. 2000. "Using Establishment Size to Measure the Impact of Title VII and Affirmative Action." *Journal of Human Resources* 35(3): 503–523.
- Chetty, R. 2006. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 90: 1879–1901.
- Coate, S., and G. C. Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83(5): 1220–1240.
- Fairlie, R. W., and W. A. Sundstrom. 1999. "The Emergence, Persistence, and Recent Widening of the Racial Unemployment Gap." *Industrial and Labor Relations Review* 52(2): 252–270.
- Fang, H., and A. Moro. 2011. "Theories of Statistical

- Discrimination and Affirmative Action.” In *Handbook of Social Economics*. Vol. 1A, J. Benhabib, M. Jackson, and A. Bisin, eds. Amsterdam: North Holland, pp. 133–200.
- Fryer, R. G., and M. O. Jackson. 2008. “A Categorical Model of Cognition and Biased Decision Making.” *B.E. Journal of Theoretical Economics* 8(1): 1–42.
- Fryer, R. G., and G. C. Loury. 2013. “Valuing Diversity.” *Journal of Political Economy* 121(4): 747–774.
- Gruber, J. 1997. “The Consumption Smoothing Benefits of Unemployment Insurance.” *American Economic Review* 87(1): 192–205.
- Holzer, H. 1998. “Why Do Small Establishments Hire Fewer Blacks Than Larger Ones?” *Journal of Human Resources* 33(4): 896–914.
- Holzer, H., and D. Neumark. 2000a. “Assessing Affirmative Action.” *Journal of Economic Literature* 38(3): 483–586.
- . 2000b. “What Does Affirmative Action Do?” *Industrial and Labor Relations Review* 53(2): 240–271.
- Ihlanfeldt, K. R., and D. L. Sjoquist. 1990. “Job Accessibility and Racial Differences in Youth Unemployment Rates.” *American Economic Review* 80: 267–276.
- Kain, J. F. 1968. “Housing Segregation, Negro Employment and Metropolitan Decentralization.” *Quarterly Journal of Economics* 82(3): 175–197.
- Kurtulus, F. A. 2011. “The Impact of Affirmative Action on the Employment of Minorities and Women over Three Decades: 1973–2003.” Unpublished manuscript.
- . 2012. “Affirmative Action and the Occupational Advancement of Minorities and Women during 1973–2003.” *Industrial Relations* 51(2): 213–246.
- Leonard, J. S. 1984. “The Impact of Affirmative Action on Employment.” *Journal of Labor Economics* 2(4): 439–463.
- Lundberg, S. J., and R. Startz. 1983. “Private Discrimination and Social Intervention in a Competitive Labor Market.” *American Economic Review* 73(3): 340–347.
- Marsden, P. V. 1994. “Selection Methods in US Establishments.” *Acta Sociologica* 37(3): 287–301.
- Mouw, T. 2000. “Job Relocation and the Racial Gap in Unemployment in Detroit and Chicago, 1980–1990: A Fixed-Effects Estimate of the Spatial Mismatch Hypothesis.” *American Sociological Review* 65(5): 730–753.
- Office of Federal Contract Compliance Programs. 2013. *Federal Contractor Compliance Manual*. Washington, DC: U.S. Department of Labor, Office of Federal Contract Compliance Programs.
- Phelps, E. S. 1972. “The Statistical Theory of Racism and Sexism.” *American Economic Review* 62(4): 659–661.
- Raphael, S. 1998. “The Spatial Mismatch Hypothesis and Black Youth Joblessness: Evidence from the San Francisco Bay Area.” *Journal of Urban Economics* 43(1): 79–111.
- Ritter, J. A., and L. J. Taylor. 2011. “Racial Disparity in Unemployment.” *Review of Economics and Statistics* 93(1): 30–42.
- Sowell, T. 2004. *Affirmative Action Around the World: An Empirical Study*. New Haven, CT: Yale University Press.
- Weinberg, B. A. 2004. “Testing the Spatial Mismatch Hypothesis Using Inter-City Variations in Industrial Composition.” *Regional Science and Urban Economics* 34(5): 505–532.
- Wilson, W. J. 1987. *The Truly Disadvantaged: The Inner City, the Underclass and Public Policy*. Chicago, IL: University of Chicago Press.
- . 1996. *When Work Disappears: The World of the New Urban Poor*. New York: Alfred A. Knopf.

Notes

1. For example, suppose an employer learns to screen from a particular group by hiring members and learning what characteristics predict productivity. Then one can view that initial set of hires as an investment to improve later screening.
2. Still, it is possible that larger establishments hire more black workers primarily because of increased legal or public pressure. In the chapter, I provide additional evidence against this alternative explanation.
3. These analyses are conducted at various levels of aggregation. For example, these correlations are measured at the individual level by Ihlanfeldt and Sjoquist (1990), at the neighborhood level by Raphael (1998), and at the metropolitan level by Weinberg (2004).
4. For example, previous researchers have used the local job density, local job growth, and the average commuting times of local workers as measures of job accessibility.
5. Alternatively, if spatial frictions are relevant, residents who find it difficult to obtain work may sort into neighborhoods with higher employment density.
6. One exception is Mouw (2000), who estimates a first differences model at the neighborhood level. His estimates remain difficult to interpret because he does not account for endogenous migration across neighborhoods.
7. For example, productive black households may move out of an MSA following widespread suburbanization, leaving less-productive black households behind.
8. Chetty (2006) generalizes the intuition behind Baily’s stylized model, demonstrating that with minor adjustments, the Baily formula holds in a more general setting that allows for a large class of realistic extensions, including arbitrary borrowing constraints, leisure benefits of unemployment, and endogenous asset accumulation or human capital investment.
9. Chetty (2006, Note 8, p. 1894).
10. As in Chetty (2006), although the model here refers to unemployment shocks the same model can be applied to other types of social insurance by relabeling the shock (e.g., injury or disability).