Essays on Labor Market Inequality

by

Conrad Miller

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

ARCHIVES
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
JUN 11 2014
LIBRARIES

© Conrad Miller, MMXIV. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature redacted

Author .....
Department of Economics
May 15, 2014

Signature redacted

Certified by .....
David Autor
Professor of Economics
Thesis Supervisor

Signature redacted

Certified by .....
Amy Finkelstein
Ford Professor of Economics
Thesis Supervisor

Signature redacted

Certified by .....
Michael Greenstone
3M Professor of Environmental Economics
Thesis Supervisor

Signature redacted

Accepted by .....
Michael Greenstone
3M Professor of Environmental Economics
Chairman, Department Committee on Graduate Theses

Essays on Labor Market Inequality

by

Conrad Miller

Submitted to the Department of Economics
on May 15, 2014, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis consists of three chapters on aspects of labor market inequality.

In chapter 1, I estimate the dynamic effects of federal affirmative action regulation, exploiting variation in the timing of regulation and deregulation across work establishments. I find that affirmative action sharply increases the black share of employees, with the share continuing to increase over time: five years after an establishment is first regulated, its black share of employees increased by an average of 0.8 percentage points. Strikingly, the black share continues to grow even after an establishment is deregulated. Building on the canonical Phelps (1972) model of statistical discrimination, I argue that this persistence is in part driven by affirmative action inducing employers to increase the precision with which they screen potential employees. I then provide supporting evidence.

In chapter 2, I study the spatial mismatch hypothesis, which proposes that job suburbanization isolates blacks from work opportunities and depresses black employment. Using synthetic panel methods and variation across metropolitan areas from 1970 to 2000, I find that for every 10% decline in the fraction of metropolitan area jobs located in the central city, black employment (earnings) declined by 1.4-2.1% (1.1-2.3%) relative to white employment (earnings). This relationship is driven primarily by job suburbanization that occurred during the 1970's. To address the potential endogeneity of suburbanization, I exploit exogenous variation in highway construction and find that highways cause job suburbanization and declines in black relative employment in a manner consistent with spatial mismatch.

In chapter 3, joint work with Isaiah Andrews, we analyze the effect of heterogeneity on the widely used analyses of Baily (1978) and Chetty (2006) for optimal social insurance. The basic Baily-Chetty formula is robust to heterogeneity along many dimensions but requires that risk aversion be homogeneous. We extend the Baily-Chetty framework to allow for arbitrary heterogeneity across agents, particularly in risk preferences. We find that heterogeneity in risk aversion affects welfare analysis through the covariance of risk aversion and consumption drops, which measures the extent to which larger risks are borne by more risk tolerant workers. Calibrations suggest that this covariance effect may be large.

Thesis Supervisor: David Autor
Title: Professor of Economics

Thesis Supervisor: Amy Finkelstein
Title: Ford Professor of Economics

Thesis Supervisor: Michael Greenstone
Title: 3M Professor of Environmental Economics

Acknowledgments

I am indebted to many for their help in completing this thesis.

I thank my advisors, David Autor, Amy Finkelstein, and Michael Greenstone. They were always generous with their time and attention, and challenged me to clarify and refine my thinking and research in ways I did not think possible. I also thank the many other faculty members who provided thoughtful feedback on my work, especially Bob Gibbons and Heidi Williams.

I thank my undergraduate advisor, Caroline Hoxby. She taught me how to think about empirical research, and showed me that I may have something to contribute. I also thank my other mentor at Stanford, Peter Blair Henry. He hired me as freshman with minimal coursework in economics, and patiently introduced me to academic research.

I thank my classmates for making the past five years enjoyable, particularly the labour group, Benjamin Feigenberg, Miikka Rokkanen, Adam Sacarny, Annalisa Scognamiglio, and Henry Swift, and honorary members Isaiah Andrews, Anil Jain, and Brad Shapiro.

I thank the National Science Foundation and Ford Foundation. Their generous funding allowed me to focus my attention on research throughout graduate school.

I thank Ron Edwards, Bliss Cartwight, and Georgianna Hawkins of the Equal Employment Opportunity Commission for facilitating access to the EEO-1 form and discrimination charge data, which I use in the first two chapters of this thesis. I thank Nathaniel Baum-Snow for providing access to the highway data used in the second chapter.

I thank Shahrzad for her love, for reminding me of life outside of graduate school, and for always believing in me, even when I had my doubts. I thank my nana, Eleanor, for her irreplaceable care packages and for reminding me to graduate sooner rather than later. Finally, I thank my parents, Cecilia and Llewellyn, for their calming presence, loving encouragement, and seemingly endless wisdom.

Contents

1 The Persistent Effect of Temporary Affirmative Action **11**

1.1 Introduction 11

1.2 Institutional Background 17

1.3 The Dynamic Effects of Affirmative Action 19

1.3.1 Data 19

1.3.2 Research Design 20

1.3.3 Main Results 24

1.3.4 Coincident Changes in Employer Characteristics 32

1.3.5 Anticipatory Behavior 36

1.3.6 Heterogeneity by Employer Size 43

1.3.7 Heterogeneity by Skill Level 45

1.3.8 Slope Fadeout 45

1.4 Causal Channels 47

1.4.1 A Screening Model with Endogenous Screening Capital 50

1.5 Screening Evidence 54

1.5.1 Data 54

1.5.2 Affirmative Action and Screening Capital 55

1.5.3 Employer Size and Black Share 56

1.6 Conclusion 64

1.A Appendix: Persistence Can Introduce Bias 66

1.B Appendix: Proofs of Propositions 67

1.C Appendix: Additional Figures and Tables 69

2 When Work Moves: Job Suburbanization and Black Employment **75**

2.1 Introduction 75

2.2 Mechanisms of Spatial Mismatch 78

2.3 Data and Implementation 80

2.3.1 Measuring Job Suburbanization 82

2.4 Empirical Strategy and Results 85

2.4.1 Baseline Estimates 88

2.4.2 Sector-Specific Suburbanization and Demand Shocks 91

2.4.3 Endogenous Migration . . . 92
2.4.4 Job versus Residential Suburbanization . . . 93
2.4.5 Highways as Exogenous Variation . . . 95
2.5 How Do Jobs Suburbanize? . . . 102
2.6 Conclusion . . . 105
2.A Appendix . . . 107

3 Optimal Social Insurance with Heterogeneity **111**
3.1 Introduction . . . 111
3.2 The Baily-Chetty Formula . . . 113
3.3 Welfare Analysis Under Heterogeneity . . . 115
3.3.1 Normalization of Utilities . . . 116
3.3.2 Stabilizing Welfare Weights . . . 116
3.3.3 Aggregating Money-Metric Utilities . . . 117
3.3.4 Alternative Approaches to Evaluating Aggregate Welfare . . . 118
3.4 Baily-Chetty Under Heterogeneity . . . 119
3.4.1 Actuarially Fair UI Under Heterogeneity . . . 119
3.4.2 UI with Uniform Taxes Under Heterogeneity . . . 120
3.4.3 Implementing the Formulas . . . 122
3.5 Calibrating the Covariance Effect . . . 122
3.6 Extensions . . . 125
3.6.1 UI with Proportional Benefits and Taxes . . . 126
3.6.2 Coefficient of Relative Prudence . . . 128
3.6.3 Sufficient Statistics for Dynamic Models . . . 130
3.7 Conclusion . . . 133

Bibliography **135**

List of Figures

1.1 Selection into Federal Contracting 24
1.2 Variation in Regulation and Deregulation Events 25
1.3 Regulation and Deregulation Event Studies 26
1.4 Regulation Event Study, by Duration 31
1.5 Regulation and Deregulation Event Studies, Within-Occupation 34
1.6 Establishment Size and Regulation and Deregulation Events 35
1.7 Deregulation Event Study, by Subsequent Growth 37
1.8 Likelihood of Future Regulation 39
1.9 Subsequent Contracts and Anticipation 42
1.10 Regulation and Deregulation Event Studies, by Skill Level 46
1.11 Fadeout Following Deregulation 48
1.12 Long Run Regulation Event Study 49
1.13 Establishment Size and Job Black Share 60
1.14 Black Share and Discrimination Charges 63
1.C.1 Summary Statistics by Event Study Year 70
1.C.2 Distribution of Contractor Spell Length 71
1.C.3 Number of Contractor Episodes by Establishment 72
1.C.4 Persistence and Bias 73

2.1 Job Suburbanization and Changes in Employment Rates, 1970-2000 90
2.2 Job versus Residential Centralization, 1970-1980 94
2.3 1947 Interstate Highway Plan 97

List of Tables

1.1 Event Study Sample, Summary Statistics 23
1.2 Regulation Event Study Estimates 28
1.3 Deregulation Event Study Estimates 30
1.4 Likelihood of Future Regulation 40
1.5 Regulation and Deregulation Event Studies, by Employer Size 44
1.6 Affirmative Action and Screening Methods 56
1.7 Employer Size and Black Share 58
1.8 Manager Race, Size, and Black Share 62
1.C.1 EEO-1 Reporting Rates by Industry, 1990 74
1.C.2 Discrimination Charge Data, Summary Statistics 74

2.1 Sample Descriptive Statistics, Cell-Level 81
2.2 Job Suburbanization By City 83
2.1 Job Suburbanization and Group and MSA Characteristics, 1970 87
2.2 Job Suburbanization and Labor Market Outcomes, By Decade, 1970-2000 88
2.3 Job Suburbanization and Labor Market Outcomes in Longer Run 89
2.4 Job Suburbanization and Labor Market Outcomes, By Group, 1970-1980 91
2.5 Job Suburbanization, Employment, and Migration, 1975-1980 93
2.6 Job versus Residential Suburbanization, 1970-1980 95
2.7 Highways and Job Suburbanization, 1970-2000 99
2.8 Highways, Residential Suburbanization, and Segregation, 1970-2000 101
2.9 Highways and Labor Market Outcomes, 1970-2000 101
2.1 Job Growth in the EEO-1 Data, 1971-1981 104
2.2 Decomposing Job Suburbanization, 1971-1981 105
2.A.1 Robustness Checks: Job Suburbanization and Labor Market Outcomes, 1970-1980 . 107
2.A.2 Job Suburbanization, Sector Composition, Labor Market Outcomes, 1970-1980 . . 108
2.A.3 Highways and Sector Growth, 1970-2000 109

3.1 Covariance Term Calibration 125

Chapter 1

The Persistent Effect of Temporary Affirmative Action

"...the [affirmative action] plan is a temporary measure, not intended to maintain racial balance, but simply to eliminate a manifest racial imbalance."

Justice William Brennan, Jr., *United Steelworkers of America v. Weber* (1979)

1.1 Introduction

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; for a review, see Fang and Moro 2011). 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 paper I study the dynamic effects of Executive Order 11246, the primary affirmative action

regulation for employment in the U.S. 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 (OFCCP 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).[1] 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). After describing my findings, I discuss how the present paper contributes to the existing body of work.

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 black share of employees, with the share continuing to increase over time. Five years after an establishment is first subject to the regulation, its black share of employees increased by an average of 0.8 percentage points. To put this magnitude in perspective, note that a 0.8 to 1.3 percentage point increase in the black share of the U.S. workforce 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 black share of employees *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).[2] In particular,

[1] Previous work finds that affirmative action regulation has had a negligible impact on female employment (Leonard 1989). Leonard posits that this may be due to the historical prioritization of minority employment in enforcement, or the secular growth of female employment over his period of study. In results not presented here, I also find that affirmative action regulation has minimal impact on female employment. By contrast, the results for Hispanic workers are qualitatively similar to those presented herein for black workers. These results for the Hispanic and female employee shares are available from the author upon request.

[2] One possible exception is Athey et al. (2000) who study how the benefits of mentoring for lower-level employees can affect optimal promotion policies. Though their focus is on promotion rather than hiring and they do not explicitly model an affirmative action intervention, the persistence found here may be reconcilable with a modified

because the policy variation exploited here varies across individual employers, it should have minimal effects on the human capital investment incentives workers 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 black share of their workforce 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.[3] 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).[4]

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' black share is increasing in employer size. While previous work documents a positive cross-sectional correlation between employer size and black share (Holzer 1998; Carrington et al. 2000), it is not clear what drives this relationship. For example, the authors of these studies posit that the relationship may

version of their model. I discuss the mentoring channel in more detail when I discuss possible mechanisms in section 1.4.

[3]While the human resources literature typically divides the hiring process into *recruitment* and *selection* or *screening* activities, I do not make this distinction. Instead, I view screening broadly as choosing the 'best' candidates from a set of potential workers.

[4]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.

be driven by workplace discrimination law, which does not cover establishments with fewer than 15 employees, or the concentration 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.

Still, it is possible that larger establishments hire more black workers due primarily to increased legal or public pressure. I provide three pieces of evidence inconsistent with this claim. First, while it is arguably firm level visibility that would affect the public scrutiny faced by any business, I find that the relationship between employer size and black share is primarily establishment level.

Second, the cross-sectional relationship between employer size and black share is *reversed* for black-run businesses: for such employers, black share is *decreasing* (and white share *increasing*) in size. If black-run employers are better able to screen black workers at baseline, then screening improvements should tend to decrease hiring rate disparities. By contrast, private employers have historically faced little legal pressure to hire more white workers.

Third, using the universe of discrimination charges filed with the Equal Employment Opportunity Commission (EEOC) from 1990-2004, I find that, conditional on size, the expected number of racial discrimination charges against an establishment is *increasing* in black share at all but the highest black share levels. Hence, for nearly all establishments, hiring more black workers does not appear to reduce dealings with discrimination law enforcement. The same is true for claims that result in what the EEOC terms 'merit resolutions'–charges with outcomes favorable to the claimant or where the EEOC finds that allegations have merit. These findings are perhaps unsurprising given that only 5% of charges claim discrimination in hiring and hence nearly all claims are filed by individuals who were employees at some point (Donahue and Siegelman 1991).

There exists a substantial literature on labor market anti-discrimination policies, including workplace discrimination law. Yet, while Title VII of the Civil Rights Act of 1964 and Executive Order 11246 were arguably two of the most controversial labor market interventions in U.S. history, we know little about their impact on the labor market (Donahue and Heckman 1991). This paper builds on an extensive literature documenting evidence that affirmative action regulation increased the black share of employees at federal contractors, at least prior to the early 1980's, when the Reagan administration significantly defunded the agency charged with the regulation's enforcement (Leonard 1984, 1990; Ashenfelter and Heckman, 1976; Goldstein and Smith, 1976; Heckman and Wolpin, 1976; Smith and Welch, 1984; Rodgers and Spriggs 1996; Kurtulus 2011).

In an important prior contribution, Kurtulus (2011) exploits within employer variation in contractor status and one of the data sets I also use here to estimate the impact of affirmative action regulation on the employment of women and minorities. Her combined estimates for black men and women imply that regulation generates an immediate 0.08 percentage point level increase in the black share of a firm's employees, that this effect is roughly unchanged two years after a firm transitions to non-contractor, and that the initial impact of the regulation *completely dissipates* as early as four years after a firm gains contractor status. Kurtulus (2011) is the first to consider the

potentially dynamic effects of the regulation on employment, an important innovation.

Building on this insight, a primary contribution of the present paper is to focus on the dynamic effects of regulation, particularly for temporarily regulated employers, and to consider the implications that persistence may have identification. I apply a flexible event study research design and find that: (1) the regulation's causal effect on black employment is substantially larger than previously estimated, particularly after Reagan-era defunding; (2) the impact accumulates over time; and (3) an establishment's black share of employees continues to *increase* many years following an employer's deregulation. I also provide evidence that the persistence I find here is not driven by anticipatory behavior, selective attrition, or measurement error.

The persistent effect of temporary regulation I document here has important implications for interpreting existing research in this literature, including Kurtulus (2011, 2012), Leonard (1984, 1990), Rodgers and Spriggs (1996), Ashenfelter and Heckman (1976), Goldstein and Smith (1976), Smith and Welch (1984), and Heckman and Wolpin (1976). Previous papers apply research designs based on comparisons of contractors to non-contractors, either across or within employers. In the presence of persistence these comparisons may substantially understate the causal impact of regulation because some employers that are currently non-contractors were previously contractors, and the minority share of those employers is still affected by the regulation. In particular, the more relevant comparison is between employers that have *ever* been contractors to those that have *never* been contractors. The flexible event study research design employed here accommodates such persistence. This may explain why in contrast to the existing literature I find that affirmative action regulation has a larger and more persistent effect on black employment.

Sociologists have argued that formal changes in personnel policies played a key role in defining both Title VII and affirmative action compliance (Dobbin 2009), though there is little evidence on whether those changes influence employee composition (Kalev et al. 2006). Using one of the data sets I use here, Holzer and Neumark (2000b) find that regulated firms search for, evaluate, and train workers more intensely. I highlight screening capital as one channel for the persistence I find here and provide novel evidence that this channel is empirically relevant.[5]

This paper most closely relates to work by McCrary (2007) and Miller and Segal (2012), both in subject and research design. They study racial hiring quotas that federal courts imposed on municipal police departments and other law enforcement agencies in the 1970's. The authors find that these quotas significantly increased black employment relative to national trends.[6] In addition, Miller and Segal (2012) find that, following the termination of these court orders, black share gains do not erode, but the share does stagnate. By contrast, I find that the black share of employees continues to *grow* even after private employers are deregulated.

There are several potential reasons why the results presented here differ from those in Miller and Segal (2012). Police departments may differ from private employers in the objective they face

[5]In the political domain, Beaman et al. (2009) find that gender quotas for leadership positions on Indian village councils lead to electoral gains for women in subsequent selections. The authors provide evidence that this persistence is in part driven by changes in voter attitudes toward female leaders.

[6]Miller and Segal (2012) also find that hiring quotas had only a marginal impact on female employment.

when choosing whom to hire. Police departments are also relatively constrained in their screening methods, often relying solely on service exams in selecting among applicants. Indeed, existing evidence suggests that departments primarily complied by adjusting their hiring threshold for the protected groups (Gaines et al. 1989). Hence, the screening investments that I argue are an important driver of persistence for private employers may not be a relevant channel for police departments.

This paper builds on a literature that investigates how employers screen workers in hiring. Autor and Scarborough (2008) show that the introduction of job testing at a large retail firm did not reduce minority hiring despite minorities performing significantly worse on the test, and generated productivity gains for both minority and non-minority hires. They argue that job testing will not decrease (and may increase) minority hiring as long as the test is unbiased relative to the preexisting screen. Relatedly, Holzer et al. (2006) and Wozniak (2012) argue that the use of criminal background checks and drug tests increases black hiring by providing information that is perceived to be more relevant for black candidates.[7] Autor (2008) argues that temporary help firms serve as a screening device for employers, pre-screening candidates and allowing employers to audition workers without the legal risks associated with firing. I argue that affirmative action induces employers to adopt new screening devices, and that a temporary regulation may generate persistent changes in screening.

Recently, researchers have provided evidence that referral hiring is an important screening device for firms (e.g. Fernandez et al. 2000; Fernandez and Weinberg 1997; Brown et al. 2012). Given that referral networks tend to display homophily (McPherson et al. 2001), this screening method is likely 'information-biased' in the sense that it provides more information about one social group of workers than another. I find that for regulated and larger employers, the most recent hire is less likely to be a referral. Relatedly, Giuliano et al. (2009) and Aslund et al. (2009) show that, all else equal, hiring managers tend to hire workers drawn disproportionately from their own race or immigrant group, respectively. Both sets of authors provide evidence that suggests these results are in part driven by managers having a comparative advantage in screening own-group workers. I argue that affirmative action causes employers to extract or more precisely interpret information about potential employees in a manner that reduces between-group employment disparities.

The remainder of the paper is organized as follows. The next section describes the relevant history and details of federal affirmative action regulation. In section 1.3 I estimate the dynamic effects of regulation and deregulation in an event study framework. In section 1.4 I discuss potential causal mechanisms and introduce the screening model. I then present supporting evidence in section 1.5. Section 2.6 concludes.

[7] Using National Longitudinal Survey of Youth data, Finlay (2009) finds evidence that after state criminal history records become available on the Internet, labor market outcomes are worse for ex-offenders. Estimates for non-offenders from highly offending groups are positive, though statistically insignificant.

1.2 Institutional Background

Issued by President Lyndon B. Johnson in September 1965, Executive Order 11246 mandates that federal contractors take 'affirmative action' to ensure nondiscrimination in their hiring and employment. While Title VII of the Civil Rights Act of 1964 outlawed discrimination on the basis of race, color, religion, sex, or national origin in all but the smallest private firms, Executive Order 11246 required that firms with federal contracts make active efforts to prevent discrimination.[8] For firms with 50 or more employees and holding $50,000 or more in federal contracts over a 12-month moving window, the requirements are more specific. In particular, such contractors are to identify underutilization of minorities and women in any occupation group relative to "availability." In identifying availability firms must consider "the availability of minorities having requisite skills in an area in which the contractor can reasonably recruit" (OFCCP 2013). Moreover, contractors are required to make "good faith" efforts to rectify underutilization, including the use of numerical goals with timetables. Broadly speaking, affirmative action mandates that federal contractors make a good faith effort to employ minorities at rates (at least) proportional to shares of local and qualified workforce, though local and qualified are not specified precisely. This regulation applies to all establishments under the firm, regardless of whether the particular facility is executing any portion of the contract. Hereafter, I will refer to Executive Order 11246 as affirmative action (AA) regulation. I will also refer to establishments as federal contractors if their parent firm meets the above size criteria.

Initially, 13 federal contracting agencies–for example, the Department of Defense and the General Services Administration–were responsible for enforcing AA regulation. Enforcement responsibilities were generally assigned on the basis of a contractor's industry irrespective of the agency contracting with the firm (Anderson 1996). The Office of Federal Contract Compliance (OFCC) was also established in the U.S. Department of Labor to advise and coordinate enforcement activities across contracting agencies. Although all agencies received guidance from the OFCC, there was wide variation across agencies in the scope and quality of their enforcement activities. In 1978, the Carter Administration consolidated the AA regulation enforcement activities under the renamed Office of Federal Contract Compliance Programs (OFCCP). Enforcement personnel from contracting agencies were reassigned to the OFCCP. In the analysis below, I use only data from 1978 onwards.

While a firm is a contractor, it is required to write an Affirmative Action Plan (AAP) for each establishment. An AAP must describe the organizational structure of the firm and establishment, identify underutilization of minorities by job group, and detail strategies, goals, and timetables for eliminating underutilization in the next year and beyond. Each AAP must be updated annually while the firm is a contractor. Contractors must also "maintain and have available records for each job on all applicants, hires, promotions, terminations, and any other selection decisions" disaggregated by minority group (OFCCP 2013).

To enforce the regulation, the OFCCP conducts compliance evaluations, reviews of a small frac-

[8]Executive Order 11246 did not cover discrimination on the basis of sex. The regulation was expanded to include women in 1967 under Executive Order 11375.

tion of covered establishments each year (about 1% of covered establishments in 2004[9]) to determine whether their AAPs are sufficient and whether they have made good faith efforts to implement their plans. These reviews focus on a contractor's performance in the last AAP year, where that calendar begins when the contractor updates their plan. The OFCCP also examines current year performance if a contractor is six or more months into its current AAP year. After the OFCCP notifies a contractor establishment that it has been selected for review, the establishment must submit its relevant AAP(s) and workforce flow data described above. Compliance evaluations consist of a desk audit and a possible site visit. As part of the desk audit, compliance officers determine whether an establishment's AAP is adequate and whether the establishment made sufficient efforts to improve minority utilization, relying on the submitted personnel data and EEO-1 form I use here and describe below. If potential violations are identified during the desk audit, the OFCCP may conduct an on-site review at the establishment. During a site visit, compliance officers further investigate potential violations, verify the firm's efforts to implement its AAP, and obtain information needed to work with the contractor to resolve any violations. Officers accomplish this in part by inspecting the contractor's facilities and reviewing its personnel files.[10]

If the OFCCP finds that a contractor is not in compliance, the OFCCP will seek a letter of commitment for minor violations or a conciliation agreement for major violations. Some of these agreements include financial settlements that involve back pay to alleged individual victims of discrimination. In 2004, the OFCCP collected $34.5 million from settlements on behalf of more than 9,000 workers. If the OFCCP and a contractor fail to resolve AA violations the OFCCP may take legal actions to penalize the contractor. The ultimate punishment for a contractor is to be debarred from doing business with the federal government, sometimes permanently. However, this outcome is quite rare. Only 43 companies were debarred up to 2001.[11] About half refused to develop an affirmative action plan or submit personnel data, while the other half did not make sufficient efforts to implement plans or violated an existing conciliation agreement. About sixty percent of debarred firms were later reinstated, and for those contractors the median period of debarment was 9.5 months (Pincus 2003).

Critically, the allocation of federal contracts is administered separately from AA enforcement. Hence, the racial composition of a firm should have no direct effect on whether it acquires a federal contract.[12] The one potential exception is very large contracts. For very large contracts, firms are formally required to 'pre-award' compliance evaluations–they must be in compliance before they

[9]There were 6,529 compliance evaluations conducted in 2004. This was actually a peak number over this period; from 2001-2010 the annual average was 4,500.

[10]There are two additional points to note about compliance evaluations. First, the targeting of early compliance evaluations appears to have been limited. A 1975 GAO report states that early compliance evaluations were primarily targeted based on employer size (GAO 1975). Leonard (1985a) confirms this. Second, at least in the early years of the regulation, the goals that employers set for themselves do not appear to act as rigid quotas. Leonard (1985b) finds that for sample of contractors in the 1970's, goals for minority share gains are positively correlated with realized gains, but the goals were rarely met.

[11]Pincus (2003) estimates that more 500,000 companies were government contractors between 1972 and 2001.

[12]Minority-owned businesses can sometimes qualify for set asides or other bid preferences for 'disadvantaged' businesses. Critically, eligibility depends on the background of the company's ownership, *not* the racial composition of its employees.

can formally initiate the contract (OFCCP 2013). In practice, very few contracts are sufficiently large to require pre-award compliance evaluations, and they are even less common for the firms I focus on below, which are not perennial contractors. Moreover, there is no requirement that an establishment be in compliance when it is not holding a federal contract.

1.3 The Dynamic Effects of Affirmative Action

1.3.1 Data

To undertake this analysis, I use establishment-level EEO-1 form data collected by the U.S. Equal Employment Opportunity Commission (EEOC) covering the years 1978-2004. Previous papers studying affirmative action regulation use versions of the same data, including Ashenfelter and Heckman (1976), Goldstein and Smith (1976), Heckman and Wolpin (1976), Leonard (1984, 1990), Smith and Welch (1984), Rodgers and Spriggs (1996), and Kurtulus (2011, 2012). As part of the Civil Rights Act of 1964, firms meeting certain size requirements are required to complete EEO-1 forms annually and submit them to the EEOC. Firms are required to report their overall racial and gender composition and the racial and gender composition of each of their establishments meeting size requirements, disaggregated by 9 major occupation groups.[13] Before 1982, all firms with 50 or more employees were required to submit EEO-1 forms. In 1982, the firm size cutoff was adjusted up to 100. For federal contractors, the cutoff was 25 employees before 1982 and 50 afterwards.[14] Firms are required to file a separate report for each establishment with at least 50 employees and the company headquarters. Establishments are consistently identified with firm and establishment identifiers. I observe each establishment's location, contractor status, and industry.[15] Moreover, over my period of study, the OFCCP primarily used the EEO-1 data to identify federal contractors.

I conduct my analysis at the establishment level. As discussed above, while regulation status is assigned at the firm level, the regulation defines compliance and is enforced at the establishment level.[16] For the analysis, I limit the sample to establishments located in metropolitan areas where the black share of workers in the data is at least 5% at some point from 1978-2004. This includes establishments from 204 metropolitan areas, where more than 80% of metropolitan area establishments are located. [17] Accordingly, this restriction does not substantively affect the results. I make the restriction to facilitate comparisons between local labor markets with significant black populations to those without, where compliance costs are presumably much smaller.

Due to the size requirements, establishments in the EEO-1 data are not representative of all

[13]The 9 occupation categories consist of: officials and managers, professionals, technicians, sales workers, administrative support workers, craft workers, operatives, laborers/helpers, and service workers.

[14]Results throughout are similar if I impose a uniform firm size cutoff of 100 employees for all establishment observations.

[15]Each of these is likely recorded with some error.

[16]This follows prior work in this literature with the exception of Kurtulus (2011, 2012), who conducts her analysis at the firm level.

[17]To define metropolitan areas, I use 1980 Census definitions.

U.S. establishments.[18] I estimate coverage rates for the EEO-1 data in 1990 in Table 1.C.1 in the Appendix. I calculate the proportion of employment accounted for in the EEO-1 data across industries by dividing EEO-1 reported employment by totals derived from County Business Patterns data for the 204 MSAs studied in this analysis. Unsurprisingly, industries that tend to have large establishments, e.g. manufacturing, are overrepresented, while industries that tend to have small establishments, e.g. services, are underrepresented. Overall, the EEO-1 data account for about 40% of total employment.

1.3.2 Research Design

I estimate the dynamic effects of AA regulation on the racial composition of regulated establishments. I focus on establishments that are *temporarily* regulated; those that are federal contractors for some initial period, but then never observed as a contractor again.

Estimating the causal effects of AA regulation is complicated by the fact that those firms subject to regulation, federal contractors, may differ from non-contractors on other dimensions that determine workforce composition, even in the absence of AA regulation. This makes simple comparisons of contractors to non-contractors difficult to interpret. Acknowledging this issue, previous researchers have estimated the causal effect of AA regulation by comparing black share *growth* across contractors and non-contractors (e.g., Leonard 1984). Comparing growth rates effectively nets out time-invariant level differences between establishments, potentially removing the influence of confounding factors from contractor to non-contractor comparisons. Still, a comparison of growth rates may be misleading if contractors and non-contractors are on different counterfactual trends. For example, firms that anticipate hiring more black workers may find it less costly to comply with AA regulation and hence may be more likely to seek federal contracts. Moreover, if the effect of AA on black share growth fades out over time–for example, if the operative channel is a constant level effect on the black share of *new hires*–then a comparison of growth rates may substantially understate the causal effect. In a further refinement of the literature, Kurtulus (2011) exploits within firm variation in contractor status, potentially alleviating concerns over selection. Previous research in this area suffers from an additional shortcoming: if regulation has an impact on employers that persists even when they are no longer contractors, previous estimates may be biased. This is because the research designs applied in existing work are based on simple comparisons of contractors to non-contractors, either within or across employers. In the presence of persistence these comparisons may substantially understate the causal impact of regulation because some employers that are currently non-contractors were previously contractors, and the minority share of those employers is still affected by the regulation.

To neutralize these concerns, I exploit variation in the timing of first and last federal contracts across establishments in an event study design. The idea is that the timing of when a firm is first or last a contractor will depend primarily on the availability and profitability of federal contracts rather than potential compliance costs, which seem unlikely to involve substantial idiosyncratic variation

[18] In addition, some firms fail to submit required EEO-1 forms.

within an employer. I estimate two sets of event studies. First, I estimate models for establishment black share that include lead and lag indicators for the first year an establishment is reported as a contractor in the data. I refer to this model as the *regulation* event study. If lead indicators are not significantly different from zero (implying that establishments that become contractors are not on pre-existing trends), I take the lag indicator coefficients as estimates of the dynamic effects of initial contractor status on establishment black share. This approach follows McCrary (2007) and Miller and Segal (2012).[19] Second, I estimate models where the event of interest is an establishment losing its status as a federal contractor, and never acquiring another contract (in the sample window). This is the next year an establishment is observed after the *last* year it is reported as contractor. About 57% of the establishments that I observe becoming contractors in the first place[20] eventually lose their contractor status. I refer to this model as the *deregulation* event study. I interpret changes in the pattern of coefficients following the event as reflecting the effects of becoming deregulated. For example, if the lag coefficients are negative or sloping downward, that would suggest fadeout of the effects of AA regulation on establishment composition.

One additional concern that applies to the research design here is that contractor status may affect establishment racial composition through channels other than AA regulation. I explore this issue further in section 1.3.4.

To identify the causal effect of AA regulation, I focus on the year an establishment is first identified as a contractor as the event of interest. I do this even though the establishment may lose and even regain contractor status later.[21] I first estimate regression models of the form:

$$Y_{it} = \alpha_i + \lambda_{d(i),t} + X_{it}\gamma + \sum_{j=a}^{b} \theta_j D_{it}^j + \epsilon_{it} \tag{1.1}$$

where Y_{it} denotes the black share of establishment i in year t, α_i and $\lambda_{d(i),t}$ are establishment and Census division by year fixed effects, X_{it} are controls for establishment size, and D_{it}^j are leads and lags for establishments first becoming contractors, defined as

$$D_{it}^j = D_i \mathbf{1}(t = \tau_i + j)$$

[19] As a test for selection bias, Kurtulus (2011) also estimates an event study for transitions from non-contractor to contractor. She finds that these transitions are associated with a slight increase in the black female share of employees on the order of 0.05 percentage points that completely dissipates in 4 years, and no detectable increase in the black male share. A key distinction between the event studies estimated in Kurtulus (2011) and those estimated here is that Kurtulus (2011) estimates a single model for all transitions from non-contractor to contractor, while I focus on the first such transition for employers that have never been contractors. This distinction is important because (a) many employers experience multiple such transitions from non-contractor to contractor and (b) regulation may have an effect that persists even when an employer is no longer a contractor. By the reasoning discussed above, in the presence of persistence, a model that treats all transitions from non-contractor to contractor the same may substantially underestimate the regulation's impact. I discuss this bias in more detail later in section 1.3.3 and in the Appendix.

[20] That is, eventual contractors that do not enter the data as a contractor.

[21] Note that I do not observe contractor status prior to an establishment's appearance in the data. This could cause me to mismeasure the event of interest. For example, an establishment may have been a contractor in a year that I do not observe it in the data, so that the first year I observe the establishment as a contractor is in fact *not* the first year the establishment was a contractor.

where D_i is an indicator for whether the establishment ever becomes a federal contractor, and τ_i is the year the establishment first becomes a contractor. I normalize the value of $\theta_{-1} = 0$. The sequence of θ_j can be interpreted as the difference in establishment black share from the year prior to first contract and j periods thereafter. For estimation, I set $a = -6$ and $b = 6$.[22]

In estimating (1.1), I exclude establishments that enter my sample as a federal contractor. I exclude these establishments from the main analysis for two reasons. First, by construction they cannot contribute directly to estimates of θ_j for $j < 0$. In other words, I cannot assess whether these establishments exhibit pre-trends because they are not observed prior to becoming a contractor. Second, I do not know the year they were first regulated, and so for any given observation I do not know how many years it has been since their regulation event. For the establishments that become contractors, I only include years of data that are in the 6-year window around the event. I do this so estimates of the event study endpoints, θ_{-6} and θ_6, are not driven by a mixture of various leads and lags. Relatedly, the sample of establishments driving identification of θ_j may vary with j because establishments are present in the data for a varying number of years. For this reason, as a robustness check I also estimate (1.1) including only non-contractors and a balanced panel of contractors. For the balanced sample, because this restriction reduces the sample size significantly, I set $a = -5$ and $b = 5$.

If AA regulation has its intended effects and the event study design is valid, we should see that θ_j is approximately 0 for $j < 0$ and positive and increasing in j for $j \geq 0$. The increase in establishment black share may be gradual rather than discontinuous because establishments are likely to adjust their racial composition by changing the composition of new hires.

I estimate an analogous event study model where the event of interest is an establishment losing its status as a federal contractor, and never acquiring another contract. I refer to this model as the *deregulation* event study. In that model, D_i is an indicator for whether the establishment ever loses its contractor status without ever regaining it (over the sample window), and τ_i is the year the establishment is last observed as a contractor. For these establishments, AA regulation is *temporary*. The results from this exercise will inform us about persistence of the regulation effect when that regulation is lifted. Note that while selection out of contractor status might be endogenous–for example, if firms exit as contractors because they experience positive shocks to compliance costs–a natural selection story would bias the results against finding persistence. In estimating this model, I only include contractors that lose their status as contractors. Hence, the comparison is between establishments that lose their status as contractors and establishments that have never been contractors. Again, for contractors and former contractors, I only include years of data that are in the 6-year window around the event. For eventual contractors that do not enter the data as contractors, I further restrict the data window to years following the regulation event. I also estimate the model including only non-contractors and a balanced panel of one-time contractors. For the balanced sample, I set $a = -5$ and $b = 5$.

I present summary statistics for the sample of establishments used in here in Table 1.1. There are

[22]The results are similar for alternative windows.

Table 1.1: Event Study Sample, Summary Statistics

	All	Non-Contractors Sample	Regulation Sample	Deregulation Sample	Overlapping Sample
Number of Establishments	569,061	161,703	63,595	85,745	36,030
Number of Firms	87,544	36,623	12,141	12,888	6,212
Establishment Size*	227	184	174	173	170
	(580)	(392)	(337)	(359)	(307)
Industry (%)					
Agricultural Services	0.3	0.4	0.2	0.2	0.2
Mining	0.9	0.4	0.4	0.3	0.2
Construction	2.2	1.6	1.4	1.5	0.9
Manufacturing	19.0	13.6	14.5	13.5	12.0
Transportation, Comm., Util.	10.2	6.1	6.1	4.4	4.3
Wholesale Trade	5.7	4.6	4.3	3.6	3.4
Retail Trade	28.3	39.2	40.8	49.4	51.0
Finance, Insurance, Real Estate	10.1	8.2	8.0	6.5	7.5
Services	23.4	26.1	24.3	20.6	20.6
Black Share Quantile*	50.0	48.3	48.4	50.3	47.8
Standardized Black Share*					
Mean	0.000	-0.003	-0.002	0.001	-0.020
Median	-0.305	-0.341	-0.336	-0.330	-0.340
Black Share of Employees* (%)	14.1	14.3	14.0	15.4	13.5
Black Share of Population, 15-64* (%)	15.5	15.5	15.2	15.6	15.0

Notes: Standard deviation in parentheses.
* Quantiles and normalizations defined at level of MSA by year cell. For regulation and overlapping sample, this is quantity at last year observed before regulation event. For deregulation sample, this is quantity at last year observed before deregulation event.

four samples of interest. There is the non-contractor sample, establishments that are never observed as contractors in the data. These establishments serve as a control group in all the event studies, helping to identify the region by year fixed effects as well as the γ coefficient. There are the full regulation and deregulation event samples. These include all eventual contractor establishments that do not enter or do not exit the data as a contractor, respectively. Finally, there is the 'overlapping' sample, which includes all eventual contractors that meet *both* criteria. To facilitate comparisons between the regulation and deregulation event studies, I focus on the overlapping sample for much of the analysis.

Non-contractors and the event study samples include establishments that are somewhat smaller than the average establishment in the data. They are more likely to be in Retail Trade. Prior to regulation, establishments in the regulation and overlapping samples have employee black shares that are very similar to non-contractors. In Figure 1.1 I plot the estimated density for the normalized establishment black share in the overlapping and non-contractor samples, where establishment black share is normalized to have mean zero and standard deviation one within MSA by year cells. For the overlapping sample, I plot the density for the normalized black share in the last year observed

Figure 1.1: Selection into Federal Contracting

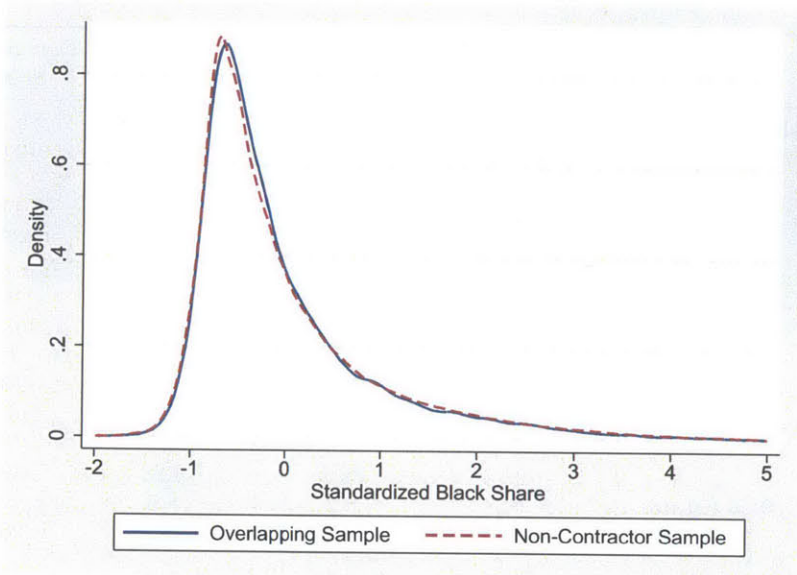

Notes: This figure is the estimated density for the normalized establishment black share in the overlapping and non-contractor samples of establishments. These samples are defined in section 1.3.2. Establishment black share is normalized to be mean zero and have standard deviation one within MSA by year cells. For the overlapping sample, this depicted density is that for the normalized black share at the last year observed before regulation event.

before regulation event. The densities are nearly identical, indicating that there is little selection into contracting based on initial black share. The average black share of employees across establishments is 14.1%, while the average black share of the working age MSA population, weighted by the number of observations, is 15.5%.

The event study design requires variation in the timing of the event of interest for the contractor sample. Figure 1.2 displays this variation for the overlapping sample. Both regulation and deregulation events vary widely in their timing.

1.3.3 Main Results

I begin with the regulation event study. In Panel A of Figure 1.3 I plot the point estimates and 95% confidence intervals of the θ_j sequence for the overlapping sample, with standard errors clustered at the firm level. The estimated model includes Census division by year fixed effects. The pattern is clear: while black share is flat for eventual contractors prior to their first contract, black share begins to increase as soon as those establishments become contractors. This suggests that the research design is valid and that the θ_j's have a causal interpretation. The effects are sizable and precisely estimated. Five years after first becoming a contractor, establishment black share grows by about 0.88 percentage points.

I assess the robustness of these estimates by altering the sample and set of controls used. The results from this exercise are presented in the top panel of Table 1.2. Columns (1) through (4) use the full regulation sample and columns (5)-(8) use the overlapping sample. Columns (4) and (8)

Figure 1.2: Variation in Regulation and Deregulation Events

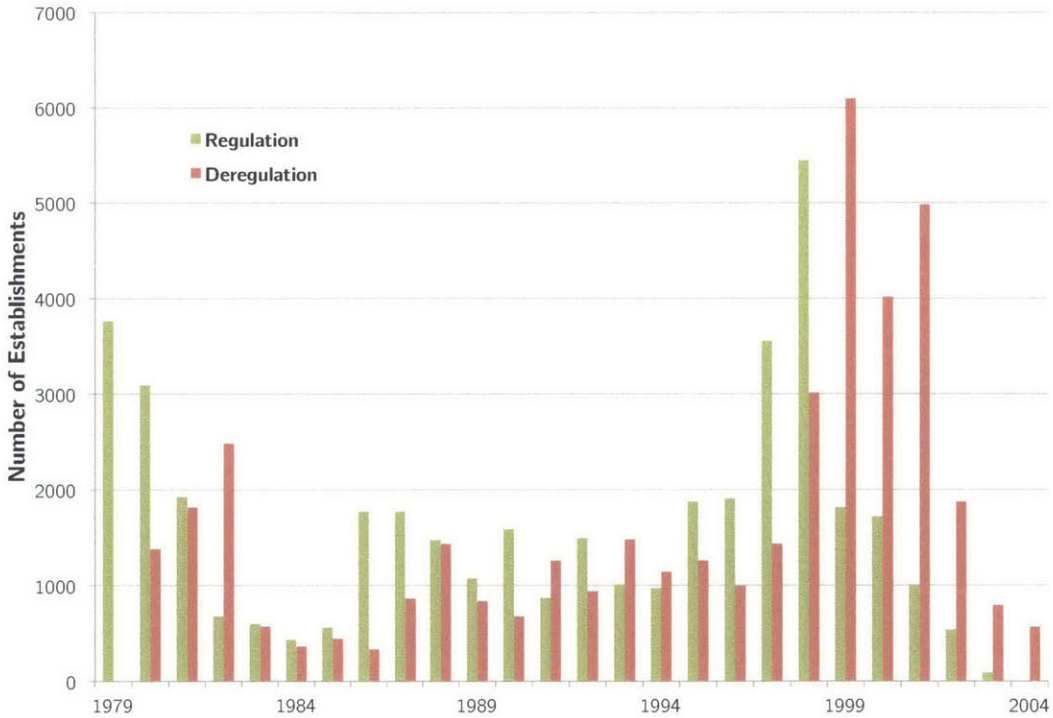

Notes: This figure is a histogram of the establishment-level regulation and deregulation events as described in section 1.3.2. Regulation events, depicted in green, are the first year an establishment is identified as a federal contractor. I exclude establishments that are contractors in the first year they are present in the data. Deregulation events, depicted in red, are the first year an establishment that was previously a contractor is: (1) not a contractor; and (2) never subsequently observed as a contractor in the data.

Figure 1.3: Regulation and Deregulation Event Studies

(A) Regulation Event

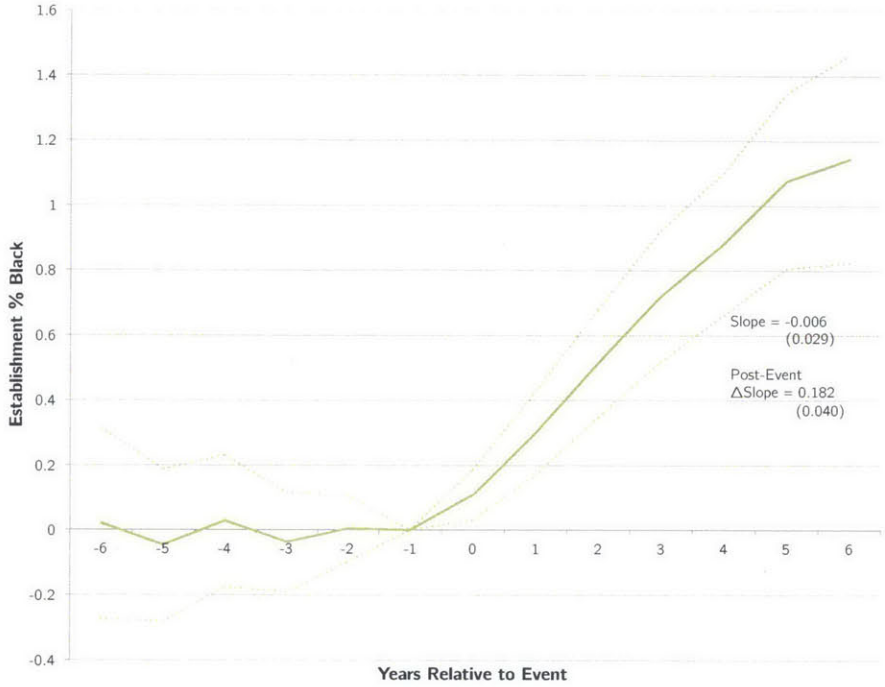

(B) Deregulation Event

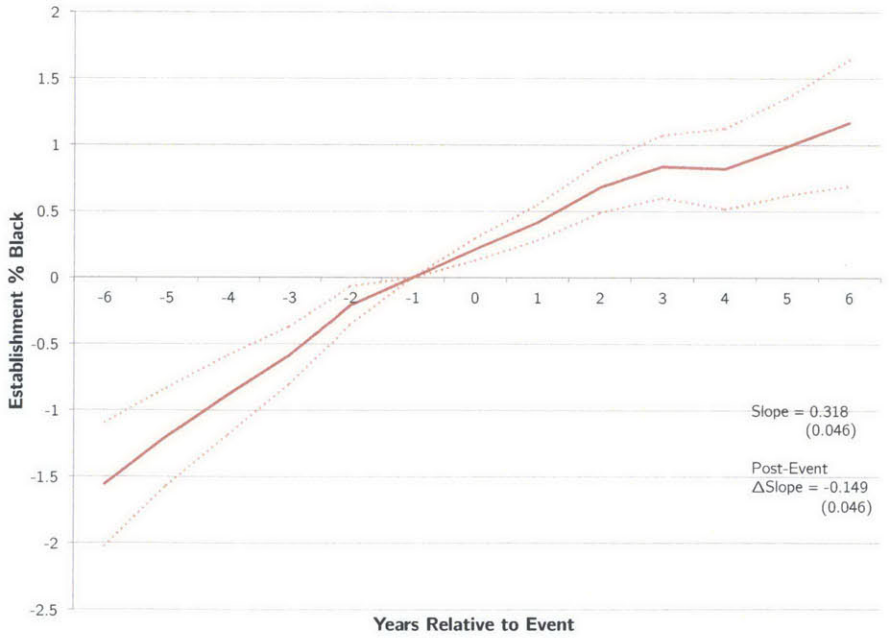

Notes: These figures plot event study coefficients and 95% confidence intervals (dotted) estimated using model (1.1) and the overlapping sample, where the outcome variable is the percent black of an establishment's employees. Panel A depicts the regulation event study; Panel B depicts the deregulation event study. The definitions of regulation and deregulation events are described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

restrict to a balanced sample of eventual contractors. Columns (1), (4), (5), and (8) include Census division by year fixed effects, columns (2) and (6) include MSA by year fixed effects, and columns (3) and (7) include Census division by 1-digit industry by year fixed effects. The coefficients from column (5) are plotted in Figure 1.3. All models include a quadratic in log establishment size. The results are similar across specifications.

Moreover, the estimated θ_j coefficients suggest a linear relationship between black share and years since first federal contract. Given this, I parameterize the event study to estimate the slope before and after the regulation event using the same sample. In particular, I estimate models of the form

$$Y_{it} = \alpha_i + \lambda_{d(i),t} + X_{it}\gamma + \beta t \times \mathbf{1}_{\exists \tau_i} + \Delta\beta(t - \tau_i + 1) \times \mathbf{1}_{(t \geq \tau_i)} + \epsilon_{it} \tag{1.2}$$

where $\mathbf{1}_{\exists \tau_i}$ is an indicator for whether an establishment is ever observed as a contractor. I present these parametric coefficient estimates in Panel B of Table 1.2. Note that, except for the balanced sample, the slope estimates put more weight on years closer to the event simply because eventual contractors are more likely to be observed in those years. The coefficient on years since first contract ranges from 0.131 to 0.182. This implies that becoming a contractor increases an establishment's black share by 0.131 to 0.182 percentage points for each year thereafter, on average. Across specifications, the pre-event slope bounces around in a small window centered at zero. The estimated slope for the balanced samples is somewhat lower, ranging from 0.132 to 0.146. Again, the pre-event slope is a relatively precise zero. Differences in the estimates may reflect differences across the samples, for example in establishment size, location, or industry.[23]

Altogether, these results suggest that AA regulation has a sizable effect on establishment personnel composition. The slope estimates are comparable to those found in Leonard (1984), though Leonard (1990) finds that AA had no impact on black employment in the 1980's using the same identification strategy.[24] I do not find this to be the case. This may be due to the differences in our research designs described above.[25]

It is important to note that many establishments included in the regulation event study are *no longer contractors* in the years following their first year as a contractor. In Figure 1.C.1 in the Appendix, I tabulate the number of eventual contractors in the overlapping sample that identify each lead and lag in Panel A, as well as the fraction of eventual contractors that are contractors in each year following the regulation event in Panel B. A year after their regulation event, only about 35% of establishments are still contractors. The same statistics are displayed for the deregulation event study.

Next, I present results for the deregulation event study. In Panel B of Figure 1.3 I plot the point estimates and 95% confidence interval of the θ_j sequence for the overlapping sample. The model

[23] For example, establishments in the balanced sample are larger and older than establishments in the full sample prior to regulation.

[24] Leonard (1984) finds that, from 1974 to 1980, affirmative action increased the relative growth rate of employment for black men and women by 0.84 and 2.13 percent annually.

[25] Note that our research strategies also require different samples. Specifically, I focus on establishments that do not enter the data as contractors. Differences in our estimates may be due in part to heterogeneous treatment effects.

Table 1.2: Regulation Event Study Estimates

	Full Regulation Sample				Overlapping Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
θ_{-6}	-0.011	-0.126	-0.271		0.020	-0.112	-0.173	
	(0.115)	(0.108)	(0.106)		(0.150)	(0.141)	(0.140)	
θ_{-5}	-0.010	-0.108	-0.221	0.075	-0.046	-0.161	-0.213	-0.031
	(0.094)	(0.089)	(0.088)	(0.111)	(0.119)	(0.110)	(0.112)	(0.125)
θ_{-4}	0.001	-0.051	-0.087	0.088	0.028	-0.036	-0.009	0.011
	(0.080)	(0.079)	(0.080)	(0.100)	(0.104)	(0.100)	(0.106)	(0.110)
θ_{-3}	-0.085	-0.124	-0.156	-0.044	-0.037	-0.100	-0.112	-0.101
	(0.060)	(0.060)	(0.062)	(0.082)	(0.078)	(0.079)	(0.082)	(0.094)
θ_{-2}	-0.016	-0.051	-0.053	-0.016	0.004	-0.044	-0.038	-0.057
	(0.041)	(0.040)	(0.041)	(0.059)	(0.052)	(0.052)	(0.055)	(0.065)
θ_0	0.095	0.098	0.092	0.072	0.109	0.109	0.079	0.087
	(0.033)	(0.033)	(0.035)	(0.066)	(0.040)	(0.040)	(0.044)	(0.079)
θ_1	0.250	0.261	0.243	0.122	0.300	0.305	0.260	0.146
	(0.053)	(0.053)	(0.053)	(0.087)	(0.066)	(0.067)	(0.067)	(0.103)
θ_2	0.435	0.458	0.433	0.325	0.516	0.529	0.465	0.368
	(0.073)	(0.071)	(0.068)	(0.102)	(0.085)	(0.084)	(0.082)	(0.113)
θ_3	0.623	0.634	0.578	0.472	0.722	0.710	0.607	0.542
	(0.093)	(0.089)	(0.081)	(0.115)	(0.104)	(0.100)	(0.092)	(0.127)
θ_4	0.795	0.814	0.780	0.576	0.883	0.874	0.783	0.654
	(0.108)	(0.104)	(0.096)	(0.129)	(0.112)	(0.108)	(0.102)	(0.145)
θ_5	0.974	1.044	1.024	0.616	1.077	1.112	1.018	0.669
	(0.127)	(0.125)	(0.115)	(0.147)	(0.138)	(0.135)	(0.123)	(0.159)
θ_6	1.047	1.172	1.154		1.144	1.241	1.162	
	(0.151)	(0.147)	(0.129)		(0.163)	(0.156)	(0.141)	
Parametric								
β	-0.005	0.014	0.037	-0.029	-0.006	0.018	0.022	-0.002
	(0.023)	(0.021)	(0.021)	(0.029)	(0.029)	(0.027)	(0.027)	(0.032)
$\Delta\beta$	0.166	0.160	0.131	0.146	0.182	0.167	0.148	0.132
	(0.034)	(0.033)	(0.031)	(0.038)	(0.040)	(0.039)	(0.037)	(0.039)
Div. × Year FEs	✓			✓	✓			✓
MSA × Year FEs		✓				✓		
Ind. × Div. × Year FEs			✓				✓	
Est. FEs	✓	✓	✓	✓	✓	✓	✓	✓
Balanced				✓				✓
# of Treated Est.		63,595		6,066		36,030		4,525

Notes: Each column reports the coefficient estimates for a regression, with standard errors in parentheses clustered at the firm level. The estimated models are regulation event studies, variants of the model (1.1) in Panel A and a parametric analog in Panel B. The definition of regulation events is described in section 1.3.2. The estimation sample includes non-contractor establishments in all columns, the regulation sample in columns (1)-(4), and the overlapping sample in columns (5)-(8). Columns (4) and (8) include only non-contractors and balanced panels of eventual contractors in the regulation and overlapping samples. All models include establishment fixed effects and a quadratic in log establishment size. Columns (1), (4), (5), and (8) include Census division by year fixed effects, columns (2) and (6) include MSA by year fixed effects, and columns (3) and (7) include Census division by 1-digit industry by year fixed effects.

includes Census division by year fixed effects. Prior to the deregulation event, an establishment's black share of employees is increasing as it is following the regulation event. Strikingly, the black share *continues to increase following deregulation.* Before the event and while regulated, an establishment's black share is increasing at a rate even larger than that found in the regulation event study. After the event, a positive slope remains. In this sense, temporary AA regulation generates ongoing increases in an establishment's black share.[26]

As for regulation event study, I assess the robustness of these estimates by varying the sample and set of controls used. I present the results in Table 1.3. Again, the estimates are comparable across specifications. I also estimate a parametric version of the model analogous to (1.2). The results are included in Panel B of Table 1.3. For the complete deregulation or overlapping samples, the pre-event slope estimates are nearly twice as large as the post-event slope estimates found for the regulation event study. After the deregulation event, this slope is half to two-thirds as large, so that the post-deregulation event and post-regulation event slopes are comparable. For the balanced sample, the pre-deregulation event slopes are about 50% as high as the corresponding post-event slope estimates for the regulation event study. There is little to no change in slope after the deregulation event. Overall, establishment black share continues to grow after the deregulation event at a rate comparable to that which emerges when establishments are first regulated.

For the overlapping subsample of eventual contractors, the regulation and deregulation event studies can be effectively combined in one plot. I estimate a series of regulation event studies, with separate estimates for eventual contractors who experience their deregulation within 1 year, 2-3 years, 4-6 years, 7-9 years, and more than 9 years. Figure 1.3.8 displays the results. In all cases, the black share of employees continues to increase even after establishments are deregulated.

The event studies can also be easily combined into one parametric regression model. Using the overlapping sample, I estimate the following model:

$$Y_{it} = \alpha_i + \lambda_{d(i),t} + X_{it}\gamma + \beta t \times \mathbf{1}_{\exists \tau_i} + \beta^R(t - \tau_i^R + 1) \times \mathbf{1}_{(t \geq \tau_i^R)} + \beta^D(t - \tau_i^D + 1) \times \mathbf{1}_{(t \geq \tau_i^D)} + \epsilon_{it} \quad (1.3)$$

where τ_i^R and τ_i^D denote regulation and deregulation event years, if applicable. I use all observation years, not restricting the data to any window around the event years. I estimate a pre-regulation slope, β, of -0.037 (with standard error 0.031); a post-regulation slope change, β^R, of 0.277 (0.059); and a post-deregulation slope change, β^D, of -0.160 (0.037).[27] The slope estimates are nearly

[26]This could be driven by temporary AA generating a persistent *level* increase in the black share of new hires, for example.

[27]With slight modification, I also estimate 1.3 excluding non-contractors from estimation. In this case, the regulation and deregulation effects are identified using only variation in the timing of events among eventual contractors. This approach is appealing in that it does not rely on non-contractors to identify the counterfactual black share for eventual contractors. However, as McCrary (2007) points out, the trend break model 1.3 is not identified using only eventual contractors. To circumvent this issue, I include observations more than 6 years prior to the regulation event, augment the model with an indicator for years more than 6 years preceding the regulation event, and limit the pre-regulation slope to apply to 6 years preceding regulation and thereafter. Reassuringly, the results are similar. I estimate a pre-regulation slope, β, of -0.009 (with standard error 0.056); a post-regulation slope change, β^R, of 0.282 (0.065); and a post-deregulation slope change, β^D, of -0.153 (0.039). The coefficient on the indicator is -0.295 with standard error 0.355, statistically insignificant at the 10% level.

Table 1.3: Deregulation Event Study Estimates

	Full Deregulation Sample				Overlapping Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
θ_{-6}	-1.312	-1.429	-1.364		-1.559	-1.678	-1.472	
	(0.192)	(0.207)	(0.061)		(0.106)	(0.253)	(0.218)	
θ_{-5}	-1.030	-1.120	-1.040	-0.775	-1.199	-1.277	-1.072	-0.799
	(0.155)	(0.166)	(0.053)	(0.111)	(0.092)	(0.197)	(0.166)	(0.327)
θ_{-4}	-0.787	-0.805	-0.726	-0.495	-0.884	-0.911	-0.802	-0.436
	(0.114)	(0.113)	(0.046)	(0.098)	(0.077)	(0.151)	(0.148)	(0.229)
θ_{-3}	-0.496	-0.520	-0.537	-0.303	-0.583	-0.633	-0.611	-0.268
	(0.086)	(0.090)	(0.036)	(0.088)	(0.063)	(0.115)	(0.113)	(0.170)
θ_{-2}	-0.225	-0.229	-0.240	-0.149	-0.209	-0.234	-0.211	-0.091
	(0.058)	(0.061)	(0.028)	(0.074)	(0.048)	(0.077)	(0.075)	(0.099)
θ_0	0.231	0.208	0.157	0.246	0.213	0.209	0.173	0.396
	(0.036)	(0.036)	(0.023)	(0.074)	(0.034)	(0.043)	(0.048)	(0.090)
θ_1	0.432	0.372	0.314	0.540	0.416	0.393	0.353	0.603
	(0.058)	(0.051)	(0.029)	(0.088)	(0.044)	(0.066)	(0.060)	(0.161)
θ_2	0.662	0.587	0.522	0.962	0.679	0.642	0.597	1.041
	(0.086)	(0.075)	(0.036)	(0.101)	(0.053)	(0.089)	(0.075)	(0.271)
θ_3	0.823	0.770	0.704	1.023	0.836	0.806	0.743	1.104
	(0.106)	(0.096)	(0.042)	(0.110)	(0.063)	(0.107)	(0.096)	(0.340)
θ_4	0.872	0.834	0.728	1.083	0.821	0.819	0.722	1.084
	(0.143)	(0.128)	(0.053)	(0.119)	(0.077)	(0.142)	(0.123)	(0.267)
θ_5	0.946	0.947	0.836	1.386	0.989	1.022	0.944	1.461
	(0.187)	(0.172)	(0.065)	(0.128)	(0.090)	(0.171)	(0.145)	(0.427)
θ_6	1.179	1.241	1.102		1.169	1.253	1.143	
	(0.233)	(0.221)	(0.081)		(0.117)	(0.231)	(0.191)	
Parametric								
β	0.274	0.288	0.268	0.203	0.318	0.334	0.290	0.221
	(0.038)	(0.040)	(0.034)	(0.059)	(0.021)	(0.048)	(0.040)	(0.088)
$\Delta\beta$	-0.101	-0.119	-0.115	0.021	-0.149	-0.161	-0.130	0.001
	(0.045)	(0.045)	(0.046)	(0.085)	(0.023)	(0.046)	(0.042)	(0.110)
Div. × Year FEs	✓			✓	✓			✓
MSA × Year FEs		✓				✓		
Ind. × Div. × Year FEs			✓				✓	
Est. FEs	✓	✓	✓	✓	✓	✓	✓	✓
Balanced				✓				✓
# of Treated Est.		85,745		5,682		36,030		2,530

Notes: Each column reports the coefficient estimates for a regression, with standard errors in parentheses clustered at the firm level. The estimated models are deregulation event studies, variants of the model (1.1) in Panel A and a parametric analog in Panel B. The definition of deregulation events is described in section 1.3.2. The estimation sample includes non-contractor establishments in all columns, the deregulation sample in columns (1)-(4), and the overlapping sample in columns (5)-(8). Columns (4) and (8) include only non-contractors and balanced panels of eventual contractors in the deregulation and overlapping samples. All models include establishment fixed effects and a quadratic in log establishment size. Columns (1), (4), (5), and (8) include Census division by year fixed effects, columns (2) and (6) include MSA by year fixed effects, and columns (3) and (7) include Census division by 1-digit industry by year fixed effects.

Figure 1.4: Regulation Event Study, by Duration

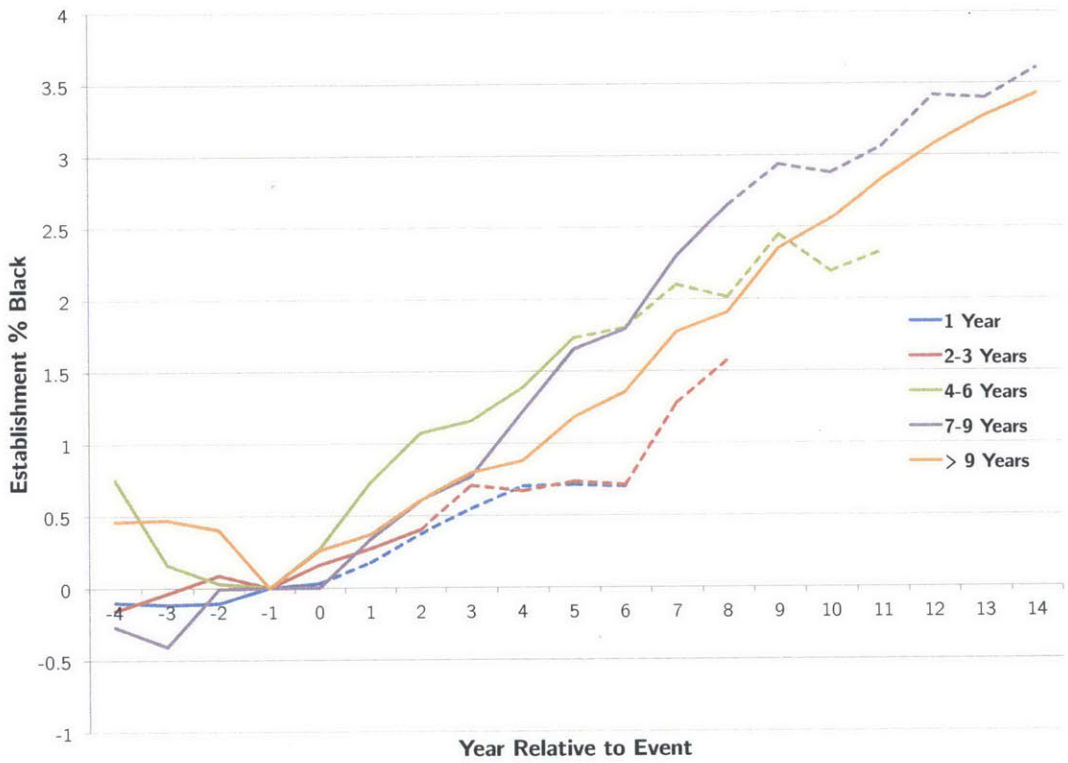

Notes: This figure plots event study coefficients and 95% confidence intervals (dotted) estimated using model 1.1 and the overlapping sample, where the outcome variable is the percent black of an establishment's employees. Each line depicts the event study estimates for a distinct group of eventual contractors, grouped by the number of years between their regulation and deregulation events. The definitions of regulation and deregulation events are described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

identical if I include a quadratic post-regulation term. I discuss the slope fadeout associated with deregulation in more detail in section 1.3.8.

The persistent effect of temporary regulation I document here has important implications for interpreting existing research in this literature, including Kurtulus (2011, 2012), Leonard (1984, 1990), Rodgers and Spriggs (1996), Ashenfelter and Heckman (1976), Goldstein and Smith (1976), Smith and Welch (1984), and Heckman and Wolpin (1976). In particular, if regulation has an impact on employers that persists even when they are no longer contractors, previous estimates may be biased. This is because the research designs applied in existing work are based on comparisons of contractors to non-contractors, either within or across employers. In the presence of persistence these comparisons may substantially understate the causal impact of regulation because some employers that are currently non-contractors were previously contractors, and the minority share of those employers is still affected by the regulation. In the Appendix, I evaluate this bias in more empirical depth.

1.3.4 Coincident Changes in Employer Characteristics

One concern with interpreting the above results is that the regulation and deregulation events involve more than changes in the set of regulations to which an establishment is subject; contractor status may have direct implications for how an establishment is organized and who it employs. I explore two potential issues directly. First, the occupational composition of an establishment may change when it becomes a contractor or loses its status as contractor. If racial composition varies systematically by occupation, an establishment's black share may change even if there are no changes in within-occupation black share. To assess the importance of occupational changes, I reestimate the event studies focusing on within-occupation changes. Second, the size or growth trajectory of establishments may change with either the regulation or deregulation events. Establishment size or growth may affect establishment composition independent of AA regulation, a question that will be explored further below. For this reason, I assess whether accounting for establishment size or growth has important implications for the above event study results.

To reestimate the event studies using within-occupation changes, I estimate (1.1) at the establishment by occupation level (rather than the establishment level.) In particular, I estimate

$$Y_{iot} = \alpha_{io} + \lambda_{d(i),t} + X_{it}\beta + \sum_{j=a}^{b} \theta_j D_{it}^j + \epsilon_{iot} \tag{1.4}$$

where o indexes occupation and Y_{not} is the black share of employees in an establishment by occupation cell. Note that (1.4) now includes occupation by establishment fixed effects. I weight observations by the current share of an establishment's workers in that occupation, and cluster standard errors at the firm level.

I plot the coefficients in Figure 1.5, alongside the analogous coefficients displayed in Figure 1.3. Panel A displays the regulation event study and Panel B displays the deregulation events study. The results are very similar to those in Figure 1.3, implying that those results are not a feature of

changes in occupational composition associated with gaining or losing contractor status.

Another employer characteristic that may vary with contractor status is establishment size. Government contracts may be sizable relative to employers' total revenue, so that employers grow or shrink when they become contractors. To explore how establishment size evolves while an establishment is regulated, I estimate separate event studies for the overlapping sample and for a subsample restricting eventual contractors to those with more than 6 years between their regulation and deregulation events. For each sample, I estimate a variant of (1.1) replacing black share with log establishment size as the dependent variable.

The results are shown in Figure 1.6. For both samples, establishment size is growing preceding the regulation event. This may be due in part to establishments anticipating a future government contract. For the restricted subsample, establishment size continues to increase following the regulation event. Over 5 years after the regulation event, establishment size increases by about 5% for this subsample.

By contrast, establishment size begins to *decrease* two years before the deregulation event and continues to decline after the event for both samples. The size responses to regulation and deregulation appear roughly symmetric, in stark contrast to how the black share of employees evolves. This exercise also serves as a falsification test for the main regulation and deregulation event study results. The broadly symmetric results for establishment size suggest that the regulation and deregulation events as constructed indeed reflect a meaningful event and its reversal.[28]

Given that establishment size decreases after the loss of contractor status, it is unclear whether the black share increases following deregulation actually reflect increases in the number of black employees. Instead, the number of black employees may stagnate or decline while the number of non-black employees declines at a faster rate. This alters the interpretation of the persistence somewhat. For example, this suggests the persistence may not come from changes in who is hired, but rather changes in relative turnover. As I will show below, as establishments grow, their black share tends to increase, while the opposite is true for declines. Hence, the size declines alone are unlikely to be responsible for the black share increases following deregulation.

To explore this question, I split the sample into establishments that decrease in size from their last year as contractor to the last year they are observed in a 6-year window following the deregulation event, and establishments that increase in size over that period. About 43% of establishments grow over this period. I then replicate the deregulation event studies separately for the two subsamples with the following modification.

To form more appropriate comparison groups for the two subsamples, it would seem natural to divide non-contractors into those that shrink and grow over comparable periods. This is complicated by the fact that there is no analogous 'event' to use as a point of reference for establishments that never become contractors. To deal with this issue, I assign pseudo 'events' to establishments that never acquire federal contracts. I do this by conditioning on two variables: the year I first observe

[28]These patterns also emerge in local labor markets with very small minority populations, suggesting the size responses are not due to regulation *per se*, but contractor status.

Figure 1.5: Regulation and Deregulation Event Studies, Within-Occupation

(A) Regulation Event

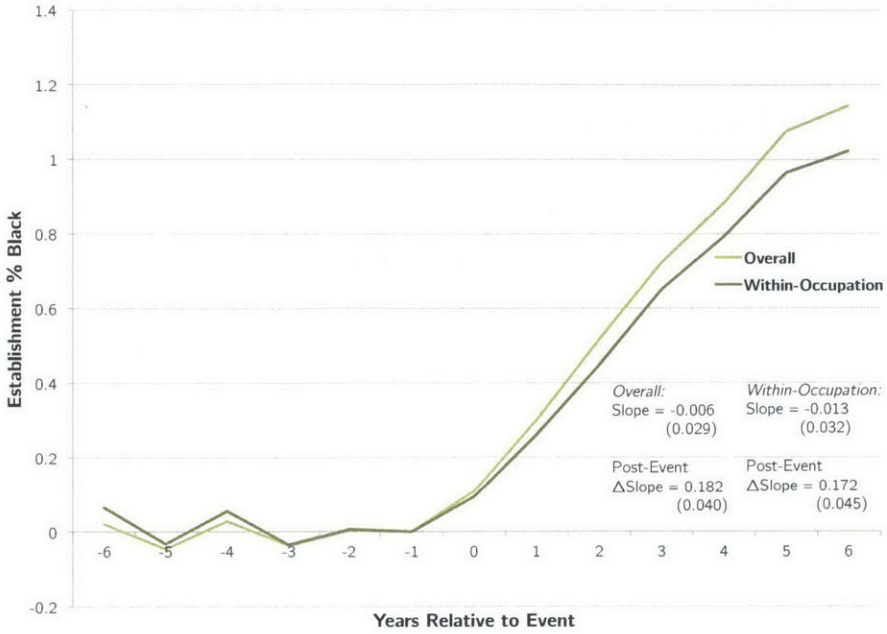

(B) Deregulation Event

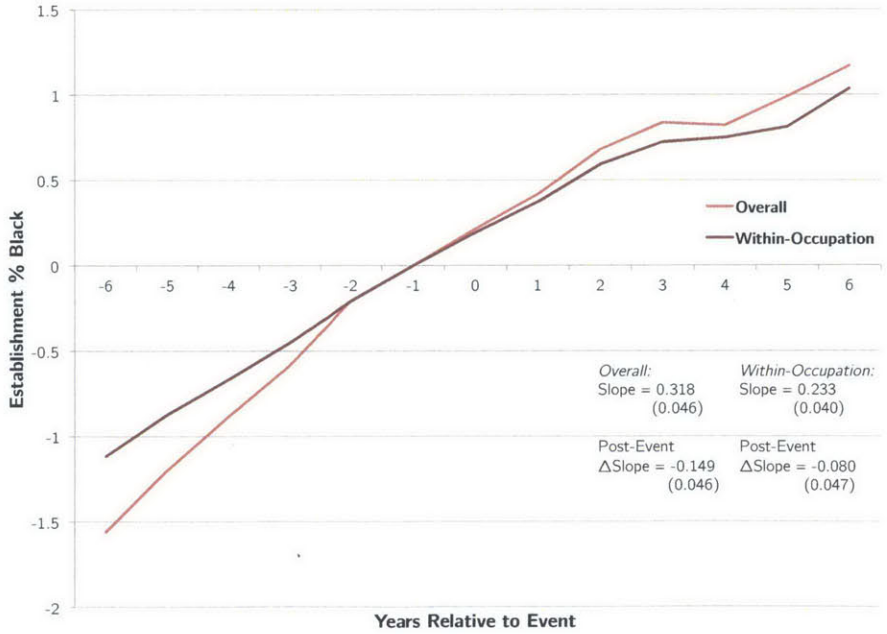

Notes: These figures plot the event study coefficients and 95% confidence intervals (dotted) estimated using model (1.1) and the overlapping sample, where the outcome variable is the percent black of employees in an establishment by occupation cell. Panel A depicts the regulation event study; Panel B depicts the deregulation event study. The definitions of regulation and deregulation events are described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level. Observations are weighted by the establishment by occupation cell's share of total establishment employment.

Figure 1.6: Establishment Size and Regulation and Deregulation Events

(A) Regulation Event

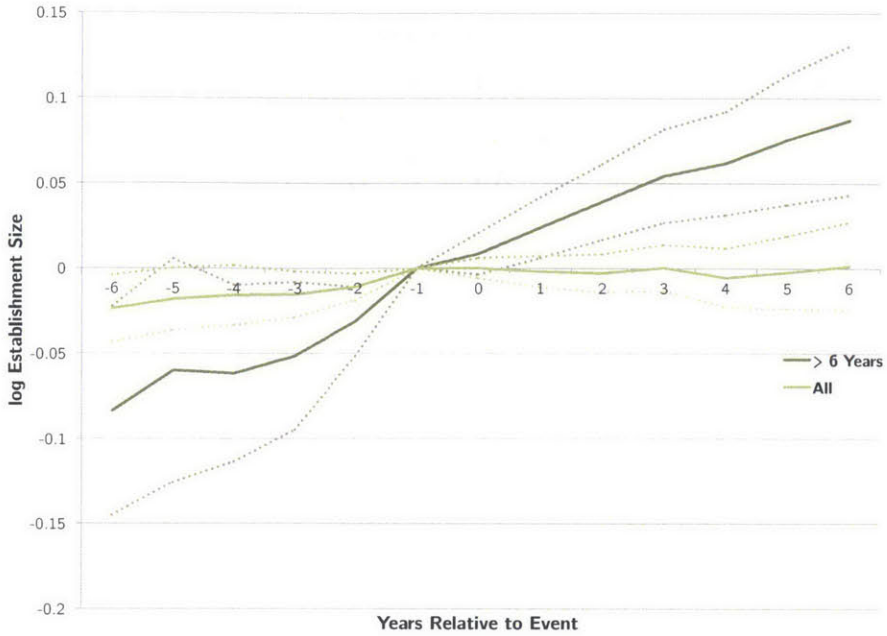

(B) Deregulation Event

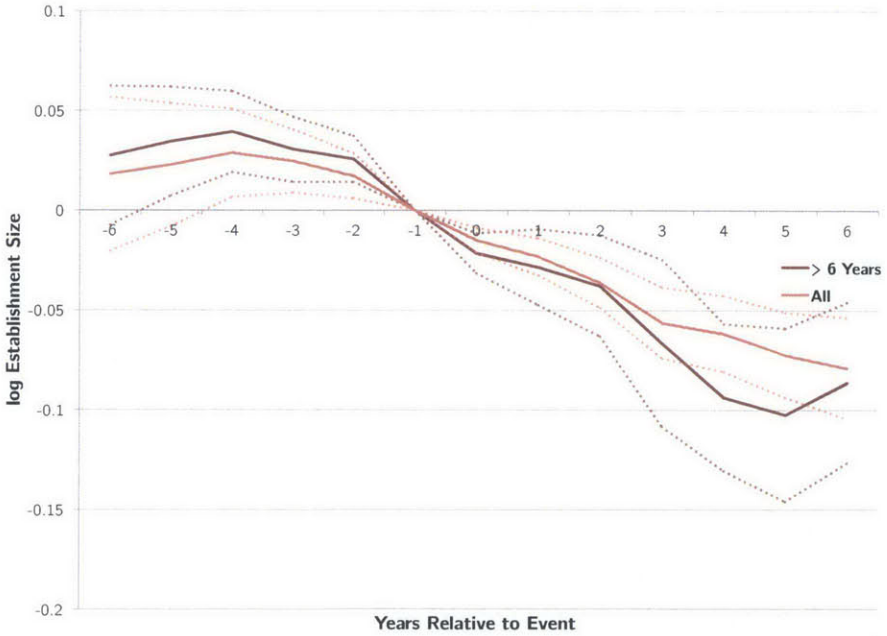

Notes: These figures plot the event study coefficients and 95% confidence intervals (dotted) estimated using model 1.1 and the overlapping sample, where the outcome variable is log establishment size. Panel A depicts the regulation event study; Panel B depicts the deregulation event study. The definitions of regulation and deregulation events are described in section 1.3.2. The '> 6 Years' line restricts eventual contractors to those with more than 6 years between their regulation and deregulation events. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects. Standard errors are clustered at the firm level.

the establishment in the data and the number of years between the first and last year. I then randomly assign an 'age' for each establishment's pseudo event, taking draws from the conditional age distribution for former contractors that lost their contractor status. Then, using this pseudo event, I similarly split the sample into establishments that decrease and increase in size following 'deregulation.' Finally, I estimate the following model separately for the two subsamples:

$$Y_{it} = \alpha_i + \lambda_{d(i),t} + X_{it}\gamma + \sum_{j=a}^{b} \theta_j D_{it}^j + \sum_{j=a}^{b} \tilde{\theta}_j \tilde{D}_{it}^j + \epsilon_{it} \tag{1.5}$$

where $\tilde{D}^j$ are analogous leads and lags for each establishments event, real or synthetic. The θ_j coefficients are the differential effects for establishments that are actually deregulated.

The results are shown in Panels A and B of Figure 1.7. In fact, the event studies for establishments that shrink and those that grow are quite similar.[29] For establishments that grow, total black employees continue to grow after the loss of contractor status.

1.3.5 Anticipatory Behavior

An additional concern with interpreting the above results is that my definition of temporary is *ex-post*. I define an establishment as temporarily subject to AA regulation if it is a contractor for some period and then never observed as a contractor again in the future. But if employers anticipate that they will become contractors again, this may blur the distinction between regulated and unregulated. Relatedly, unregulated employers may perceive that increasing their black share will improve their chances of acquiring a future contract. In this section, I address these two concerns.

First note that it is not *a priori* clear how the regulation would incentivize anticipatory behavior in personnel decisions. In particular, the regulation does not require that firms be in compliance when they are not regulated. Moreover, as discussed in section 1.2, a firm's racial composition should have no direct bearing on whether it is able to acquire a future contract. Nevertheless, it is possible that employers do not understand this, that contract allocation depends on racial composition for reasons outside of the law, or that adjustment costs make anticipatory behavior sensible. I assess the potential importance of anticipatory behavior in driving the previous results in two ways. First, I estimate how the likelihood of future regulation evolves after an employer loses its contractor status. Second, I test whether the pattern of persistence following deregulation differs in cases when *deregulation* is *temporary*–that is, transitions from contractor to non-contractor where an establishment *does* win a subsequent contract. Third, I investigate whether current black share predicts future contractor status.

For the first exercise, I compute the conditional probability for a firm to acquire a future contract as a function of the number of years since the firm last held a contract. This enables two types of comparisons. First, I compare the acquisition likelihood for establishments that previously held a

[29] In results not shown here, I find that the results are not sensitive to whether establishment size is included as a control variable

Figure 1.7: Deregulation Event Study, by Subsequent Growth

(A) Shrinking Establishments

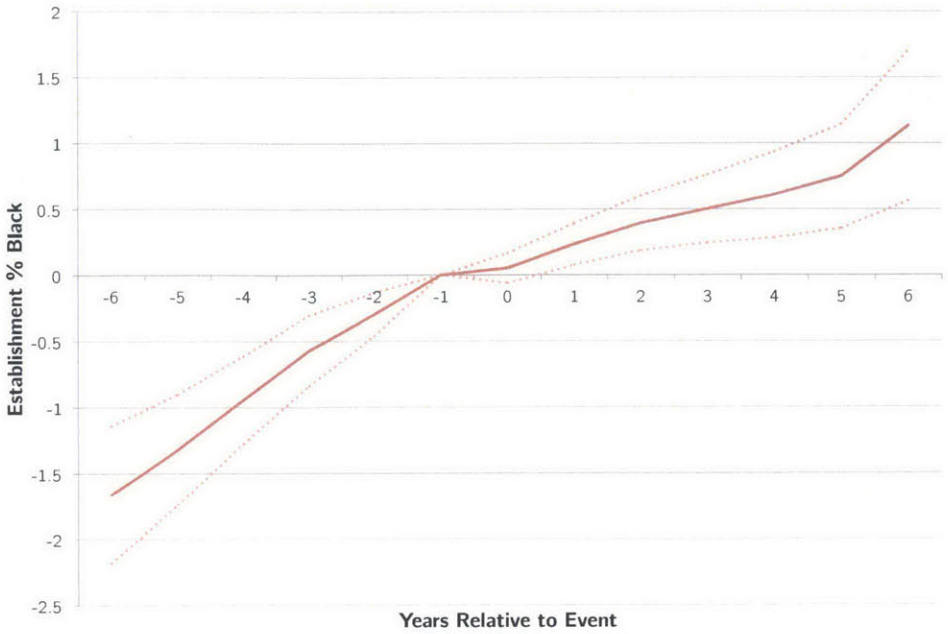

(B) Growing Establishments

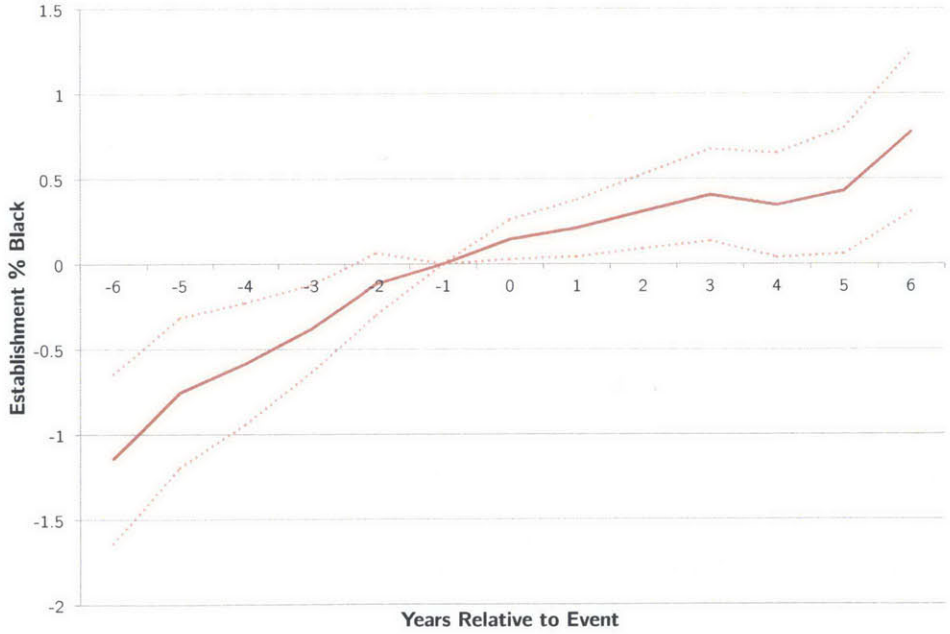

Notes: These figures plot the deregulation event study coefficients and 95% confidence intervals for various outcome variables. The definition of deregulation events is described in section 1.3.2. Panels A and B plot estimates of model (1.5) using only establishments in the overlapping sample that shrink and grow following the deregulation event. See section 1.3.4 for details. Pseudo event years are assigned to non-contractors as described in section 1.3.4. The outcome variable for these two panels is the percent black of employees. In all models the coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

contract to those that have never held a contract. Second, I explore how the acquisition likelihood evolves with the number of years since a firm last held a contract. In the presence of important anticipatory behavior, we would expect to see a relationship between the likelihood of future contract acquisition and contemporaneous minority share gains. To make the findings applicable to the main results using the overlapping sample, I limit the analysis to establishments that do not enter the data as federal contractors.[30]

These acquisition likelihoods are plotted in Figure 1.8. The vertical axis denotes the fraction of firms that acquire a federal contract in the future, either in the next year, as in Panel A, or in the next three years, as in Panel B. To avoid censoring, I exclude observations from 2004 and 2002-2004 in constructing Panel A and Panel B. The purple line depicts this likelihood for establishments that have previously held a federal contract, but have not have held a contract for a given number of years, as marked on the horizontal axis. For the 0 value on the horizontal axis, the purple line denotes the likelihood for current contractors. The blue line depicts the likelihood for establishments that have never held a federal contract, and the horizontal axis marks the number of years they've been observed in the data.

I estimate regression versions of these plots with controls that include a quadratic in establishment size, 1-digit industry fixed effects, and Census division by year effects to adjust for regional and temporal variation. In the regression models, I also examine acquisition likelihoods for a five year window and whether an establishment ever becomes a contractor as observed in the data. For these two likelihoods I exclude observations from 2000-2004 and 2004, respectively. I also try limiting the former contractors to those who have been previously observed as contractors for at least three years. I present the results in Table 1.4.

There are three points to note from Figure 1.8 and Table 1.4. First, the likelihood that a firm that has never held a contract acquires one in the future is roughly constant in the age of the firm, declining slightly with age.

Second, after a year without a contract, one-time contractors are only slightly more likely to acquire a future contract as firms that have never held a contract. While current contractors are about 45 percentage points more likely to hold a contract in 1 or 3 years than firms that have never held a contract, this difference reduces to 5 and 13 percentage points after one year without a contract. After 4 years, the difference reduces further to 0 and 5 percentage points. Yet, Figure 1.3 shows that such firms continue to increase their minority share after going four years without a contract, while there is no evidence of anticipatory behavior for firms that have never held a contract.

Third, for one-time contractors, there is a substantial decline in the likelihood of acquiring a future contract with each passing year without a contract, especially after the first year. While about 55% of firms who hold a contract in one year are contractors in the next, after one year without a contract, only about 25% of former contractors hold one in the next. After 4 years, this likelihood declines to about 8%. By contrast, robust black share increases continue following

[30]The pattern of results is similar if I do not make this restriction.

Figure 1.8: Likelihood of Future Regulation

(A) Likelihood of Regulation in Next Year

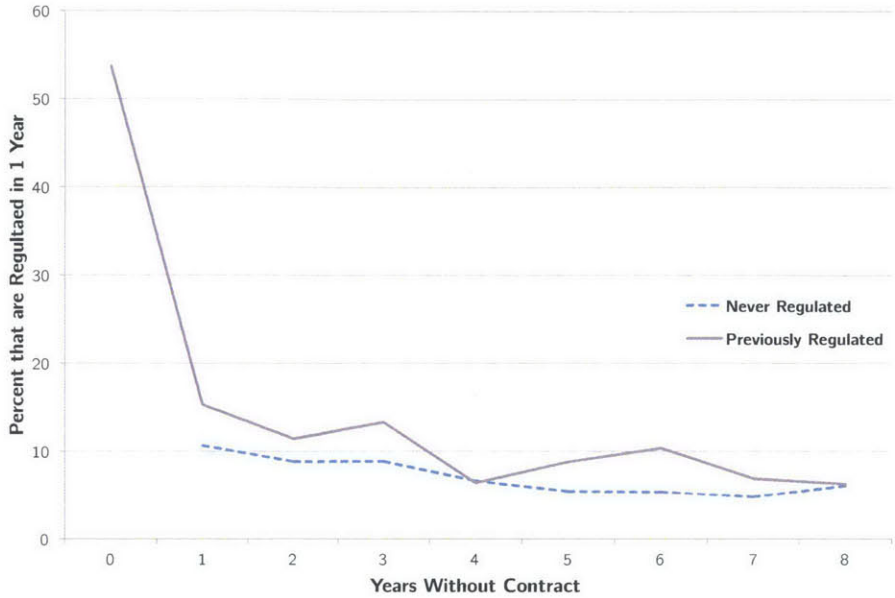

(B) Likelihood of Regulation in Next 3 Years

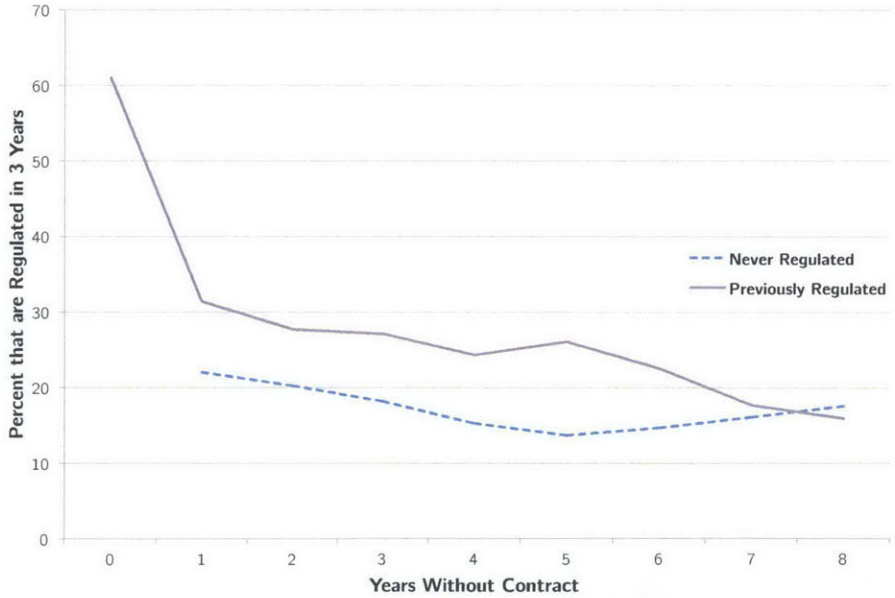

Notes: This figure plots the probability that an establishment acquires a future contract in some time period, conditional on the number of years since the establishment last held a contract or first appeared in the data. The vertical axis denotes the fraction of firms that acquire a federal contract in a given time period, either in the next year, as in Panel A, or in the next three years, as in Panel B. Panel A includes data from 1979 to 2003. Panel B includes data from 1979 to 2001. The purple line depicts this likelihood for firms that have previously held a federal contract, but have not have held a contract for a given number of years, as marked on the horizontal axis. For the 0 value on the horizontal axis, the purple line denotes the fraction of current contractors that will be contractors in the given time period. The blue line depicts this likelihood for establishments that have never held a federal contract, and for these establishments the horizontal axis marks the number of years they've been observed in the data.

Table 1.4: Likelihood of Future Regulation

	All				Experienced			
	Acquires Contract in Next							
	1 Year (1)	3 Years (2)	5 Years (3)	Ever (4)	1 Year (5)	3 Years (6)	5 Years (7)	Ever (8)
Year 1 × Former Contractor	0.048	0.114	0.138	0.127	0.084	0.181	0.218	0.166
	(0.014)	(0.015)	(0.016)	(0.012)	(0.027)	(0.023)	(0.028)	(0.018)
Year 2 × Former Contractor	0.042	0.090	0.133	0.088	0.068	0.163	0.196	0.116
	(0.012)	(0.020)	(0.021)	(0.014)	(0.019)	(0.041)	(0.044)	(0.021)
Year 3 × Former Contractor	0.054	0.072	0.110	0.057	0.095	0.130	0.181	0.065
	(0.016)	(0.017)	(0.023)	(0.013)	(0.025)	(0.038)	(0.041)	(0.019)
Year 4 × Former Contractor	-0.005	0.040	0.047	0.018	-0.003	0.016	0.064	0.002
	(0.000)	(0.019)	(0.024)	(0.014)	(0.014)	(0.040)	(0.050)	(0.026)
Year 5 × Former Contractor	0.023	0.054	0.048	0.022	0.028	0.070	0.059	0.003
	(0.016)	(0.026)	(0.027)	(0.017)	(0.039)	(0.060)	(0.060)	(0.035)
Year 6 × Former Contractor	0.031	0.048	0.041	0.039	-0.018	0.055	0.049	0.008
	(0.017)	(0.024)	(0.025)	(0.018)	(0.019)	(0.052)	(0.061)	(0.038)
Year 7 × Former Contractor	0.010	0.033	0.026	0.029	0.061	0.084	0.086	0.037
	(0.016)	(0.022)	(0.024)	(0.020)	(0.048)	(0.056)	(0.070)	(0.050)
Year 8 × Former Contractor	0.020	0.022	0.060	0.032	0.043	0.043	0.059	-0.002
	(0.015)	(0.020)	(0.027)	(0.020)	(0.034)	(0.045)	(0.062)	(0.037)
Year 2	-0.020	-0.011	-0.008	0.009	-0.020	-0.011	-0.007	0.009
	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.004)
Year 3	-0.016	-0.014	-0.009	0.014	-0.016	-0.014	-0.009	0.015
	(0.005)	(0.006)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)	(0.005)
Year 4	-0.023	-0.030	-0.023	0.010	-0.024	-0.030	-0.023	0.011
	(0.005)	(0.007)	(0.007)	(0.006)	(0.005)	(0.007)	(0.007)	(0.006)
Year 5	-0.035	-0.047	-0.031	0.004	-0.034	-0.045	-0.028	0.008
	(0.005)	(0.008)	(0.009)	(0.007)	(0.005)	(0.008)	(0.009)	(0.007)
Year 6	-0.039	-0.068	-0.053	-0.017	-0.039	-0.065	-0.049	-0.013
	(0.005)	(0.008)	(0.011)	(0.008)	(0.005)	(0.008)	(0.011)	(0.008)
Year 7	-0.051	-0.085	-0.068	-0.026	-0.049	-0.078	-0.064	-0.022
	(0.005)	(0.010)	(0.013)	(0.009)	(0.005)	(0.09)	(0.013)	(0.010)
Year 8	-0.061	-0.083	-0.073	-0.024	-0.060	-0.079	-0.073	-0.022
	(0.007)	(0.011)	(0.013)	(0.010)	(0.006)	(0.011)	(0.014)	(0.010)
Division by Year FEs	✓	✓	✓	✓	✓	✓	✓	✓
N Observations	836,625	736,595	646,244	836,625	698,967	616,620	546,132	698,967

Notes: Each column reports the coefficient estimates for a regression, with standard errors in parentheses clustered at the firm level. Columns (1), (4), (5), and (8) include data from 1979 to 2003. Columns (2) and (7) include data from 1979 to 2001. Columns (3) and (7) include data from 1979 to 1999.

deregulation.

These results suggest that persistence shown in Figure 1.3 is unlikely to be driven by firms anticipating that they will become regulated again.[31]

If deregulated establishments continue to increase their black share because they anticipate becoming contractors again, we may expect this anticipatory behavior to be more pronounced the higher the likelihood of a future contract. One could test for anticipatory behavior by comparing the persistence for establishments that believe they are more likely to become a contractor in the future to the persistence for establishments that believe a subsequent contract is less likely. Unfortunately, establishment expectations over future contractor status is not observable. Instead, I compare the behavior of establishments following a deregulation event–transitions from contractor to non-contractor where the establishment never becomes a contractor again as observed in the data–to behavior following an analogously defined 'temporary deregulation' event–transitions from contractor to non-contractor where an establishment *does* win a subsequent contract. In particular, I focus on an establishment's *last* temporary deregulation event observed in the data.The motivating assumption is that, on average, in cases when establishments *do* win a subsequent contract *ex-post*, they believed the likelihood of winning a contractor was higher *ex-ante* than in cases when establishments *do not* win a subsequent contract. If evolution of establishment black share is similar following both types of events, this would suggest that the observed persistence is not driven by anticipatory behavior.

I estimate and compare event study models for the deregulation event and the temporary deregulation event. For the former, I use the overlapping sample as in 1.3. I estimate the temporary deregulation event study model using non-contractors and eventual contractors that do not enter the data as contractors and experience a temporary deregulation event. As above, for the contractors and former contractors, I only include years of data that are both (1) in a 6-year window around the event of interest and (2) after the establishment's regulation event. For the temporary deregulation event study, I further restrict the data to years prior to the establishment's subsequent contract. Note that the same eventual contractor can serve in the estimation of both event study models.

I present the results for this exercise in Figure 1.9. The light red line depicts the same deregulation event study estimates for the overlapping sample that are shown in 1.3. The dark red line depicts analogous event study estimates where the event of interest is *temporary deregulation*, described above. The patterns are strikingly similar. In either case, the black share of employees continues to increase following deregulation at a similar slope. This suggests that the observed persistence is not driven by anticipatory behavior.

Firms may perceive that increasing the black share of their employees will improve their chances of acquiring a future contract. If so, AA regulation may alter the personnel decisions of *unregulated* employers. This would complicate the interpretation of temporary AA regulation in this context.

Two pieces of evidence suggest a firm increasing its black share does not make the acquisition of

[31]Still, if firms face substantial adjustment costs in reversing any changes in personnel practices made in response to AA regulation, former contractors may set a lower threshold for anticipatory behavior than firms who have never held a contract.

Figure 1.9: Subsequent Contracts and Anticipation

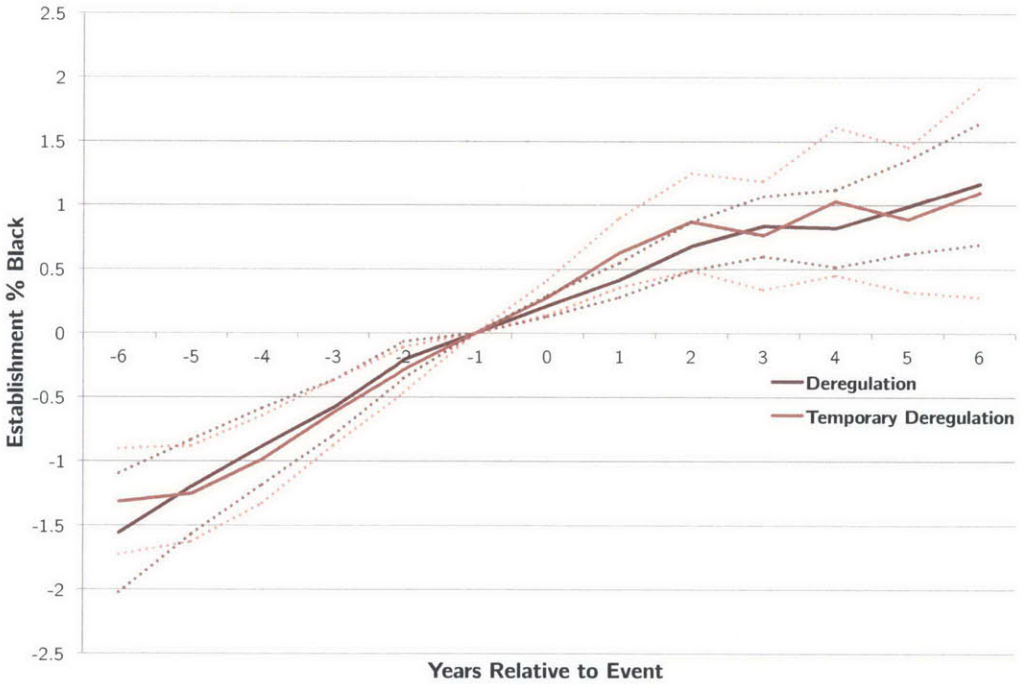

Notes: This figure plots event study coefficients and 95% confidence intervals (dotted) estimated using model (1.1), where the outcome variable is the percent black of an establishment's employees. The light red line depicts the deregulation event study estimates using the overlapping sample and the standard definition of the deregulation event, as described in section 1.3.2. The dark red line depicts event study estimates where the event of interest is 'temporary deregulation': a transition from contractor to non-contractor, where an establishment eventually regains its contractor status. The estimation sample is limited to non-contractors and eventual contractors that do not enter the data as contractors and experience a 'temporary deregulation' event. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

future contracts more likely. First, as discussed in section 1.2, the OFCCP has little to no explicit influence on contract allocation. Second, before eventual contractors are first regulated, their black share is very similar to those that never become contractors. This is evident in Figure 1.1. Moreover, initial black share is slightly *negatively* correlated with future contractor tenure.

1.3.6 Heterogeneity by Employer Size

In this section, I explore how the regulation's effect varies with employer size. This exercise serves two purposes. First, I assess whether estimates are sensitive to the selective attrition of establishments from the data. Second, I exploit the fact that compliance evaluations are targeted based on employer size (Leonard 1985a) to examine whether the regulation's impact is more substantial where enforcement is stronger.

The absence of pre-existing trends in the regulation event study suggests that the black share gains following AA regulation indeed reflect the causal effect. However, the event study may produce biased estimates for the causal effect if establishments selectively exit from the data. The size thresholds for who must submit EEO-1 forms magnifies this concern. In particular, some firms that are near the threshold may alter their size to avoid reporting requirements.

To assess the potential role of selective attrition in producing the above results, I re-estimate both the regulation and deregulation event studies restricting estimation to establishments whose parent firms have at least 150 employees prior to their first federal contract. For eventual contractor establishments, I base this restriction on firm size in the latest year an establishment is observed prior to their regulation event. For non-contractor establishments, I use firm size in the latest year an establishment is observed prior to their *pseudo* regulation event, where pseudo event events are randomly assigned as described in section 1.3.4, based on the year I first observe the establishment in the data and the number of years between the first and last year. These establishments are not near the size threshold, and so any manipulation to avoid reporting seems unlikely. Note that over 90% of establishments in the overlapping sample satisfy this criteria.

The results are shown in column (2) of Table 1.5. Panel A presents regulation event slope estimates, while Panel B presents deregulation event slope estimates. The estimates for both event studies are very similar to those using the full overlapping sample. I conclude that selective attrition is unlikely to be an important concern here.

Leonard (1985a) studies the targeting of compliance evaluations conducted by the Department of Defense over the late 1970's. He finds that contractor establishments that are part of multi-establishment firms are substantially more likely than singleton contractor establishments to be subject to a compliance evaluation. The likelihood of review is generally increasing in establishment size, though the relationship is concave. Motivated by these findings, I explore how the response to regulation depends of whether an establishment is part of a larger company, and establishment size. Note that I focus on a later period than Leonard (1985a), and the targeting of compliance evaluations has likely changed over time. Hence, these results should be interpreted with caution.

The results are presented in columns (3)-(6) in Table 1.5. Columns (3) and (4) report estimates

Table 1.5: Regulation and Deregulation Event Studies, by Employer Size

	All	Firm Size $\geq$ 150	Single Establishment	Multi Establishment	Small Establishment	Large Establishment
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Regulation Event						
β	-0.006	-0.008	-0.027	-0.022	0.002	-0.017
	(0.029)	(0.031)	(0.032)	(0.030)	(0.042)	(0.027)
$\Delta\beta$	0.182	0.190	0.017	0.202	0.170	0.192
	(0.040)	(0.042)	(0.037)	(0.039)	(0.054)	(0.050)
Panel B: Deregulation Event						
β	0.318	0.314	-0.030	0.314	0.367	0.266
	(0.046)	(0.046)	(0.041)	(0.047)	(0.058)	(0.056)
$\Delta\beta$	-0.149	-0.141	0.014	-0.135	-0.227	-0.074
	(0.046)	(0.048)	(0.046)	(0.050)	(0.054)	(0.061)
Div. $\times$ Year FEs	✓	✓	✓	✓	✓	✓
Est. FEs	✓	✓	✓	✓	✓	✓
# of Treated Est.	36,030	33,314	4,703	31,327	18,519	17,511

Notes: Each column reports the coefficient estimates for a regression, with standard errors in parentheses clustered at the firm level. The estimated models are parametric regulation and deregulation event studies in Panel A and Panel B. The definition of regulation and deregulation events is described in section 1.3.2. The estimation sample in column (1) includes non-contractor establishments and the overlapping sample. Column (2) restricts the analysis to establishments that are part of firms with at least 150 employees. Column (3) restricts the analysis to singleton establishments, and column (4) restricts the analysis to establishments that are part of multi-establishment firms. Column (5) restricts the analysis to establishments with less than 100 employees, and column (6) restricts the analysis to establishments with at least 100 employees. All columns include Census division by year fixed effects. For eventual contractor establishments, these restrictions are based on the latest year an establishment is observed prior to their regulation event. For non-contractor establishments, these restrictions are based on the latest year an establishment is observed prior to their *pseudo* regulation event, where pseudo event events are assigned based on the year I first observe the establishment in the data and the number of years between the first and last year. More details on how pseudo events are assigned are described in section 1.3.4.

based only on singleton establishments and establishments that part of multi-establishment firms, respectively. Column (5) reports estimates based only on establishments with fewer than 100 employees, while column (6) reports estimates based on establishments with 100 or more employees. The regulation appears to have little effect on singleton establishments. Larger establishments experience larger black share gains following the regulation event, and larger gains following deregulation, though similar patterns emerge for smaller establishments. Overall, it appears that establishments that are more likely to be evaluated by regulations respond more to regulation.

Note that while only about 13% of establishments in the overlapping sample are singletons, they represent over 70% of sample firms. The significant heterogeneity found here implies that an analysis that weights firms equally, as in Kurtulus (2011; 2012), rather than establishments, as done in the present paper and previous work in the literature, will produce estimates of the regulation's impact that are substantially smaller in magnitude.[32]

1.3.7 Heterogeneity by Skill Level

While AA regulation generates a sharp increase in minority share growth, and most of these gains are within-occupation, it is unclear what kinds of jobs are driving this growth. To clarify this, I repeat the within-occupation event study analyses separately by occupation skill level. Following Acemoglu and Autor (2011), I divide the occupation groups defined in the EEO-1 data into three skill groups: high, middle, and low. I label officials and managers, professionals, and technicians as "high skill"; sales workers, administrative support workers, craft workers, and operatives as "middle skill"; and laborers/helpers and service workers as "low skill."

I present the results in Figure 1.10. Event study patterns are similar across occupation groups. In absolute terms, the effect on black share is largest in middle skill occupations. Five years after the regulation event, the black share of employees in middle skill occupations increased by 0.8 percentage points. Estimates for high and low skill occupations are similar in magnitude at about 0.6 percentage points, though they are relatively imprecise. In the EEO-1 data, 7.0%, 12.6% and 20.5% of high skill, middle skill, and low skill workers are black. Hence, in proportional terms, the effects of AA are similar for both high and middle skill occupations.

1.3.8 Slope Fadeout

The deregulation event study results suggest that while the black share of employees continues to increase following deregulation, this persistence is not complete. For example, in Panel B of Figure 1.3, the post-deregulation event slope is about 35% smaller than the pre-event slope. Moreover, Figure suggests the degree of persistence may depend on an establishment's experience as a contractor. In this section, I explore possible fadeout in more detail.

Though deregulation is associated with a slope decrease in black share gains, it is not clear whether this decrease is due to deregulation *per se* or if this decrease would occur in the absence of

[32] Weighting by establishment size yields results similar to those presented here.

Figure 1.10: Regulation and Deregulation Event Studies, by Skill Level

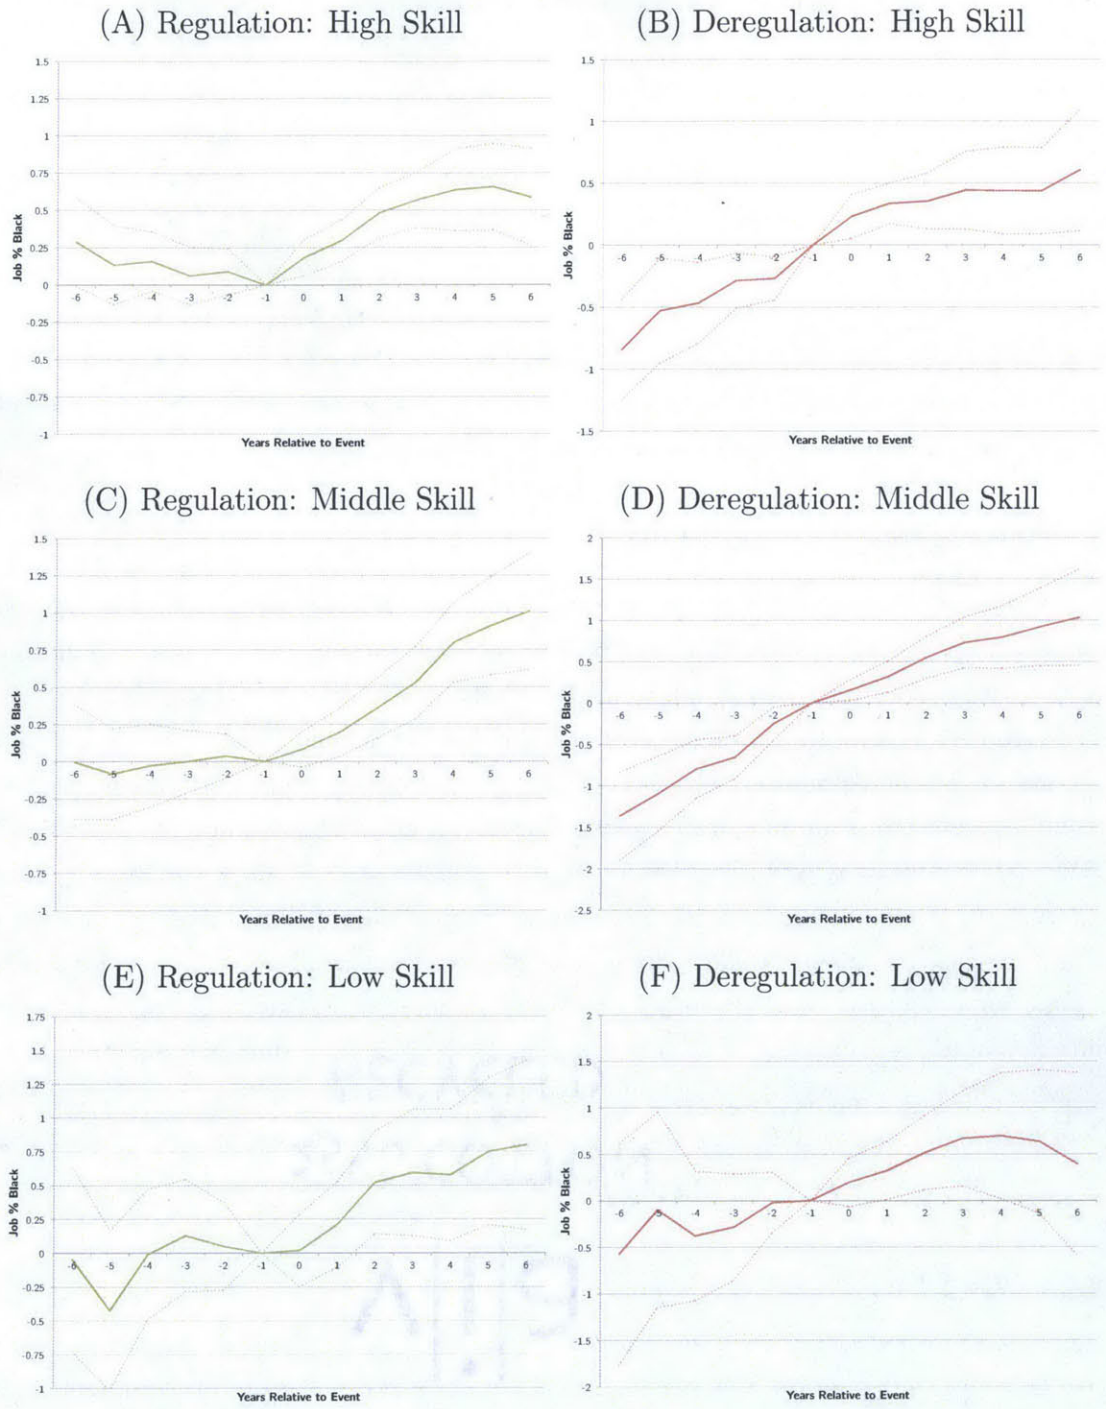

Notes: These figures plot event study coefficients and 95% confidence intervals (dotted) estimated using model 1.4 and the overlapping sample, where the outcome variable is the percent black of employees in an establishment by occupation cell. The event studies are estimated separately for high skill (managers, professionals, technicians), medium skill (sales workers, administrative support, craft workers, operatives), and low skill (laborers and service workers) occupations. The left column of panels depicts regulation event studies; right column depicts deregulation event studies. The definitions of regulation and deregulation events are described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level. Observations are weighted by the establishment by occupation cell's share of total establishment employment in the corresponding skill group.

deregulation. The slope must decrease at some point given that the black share is bounded above, where the ceiling depends on the availability of black workers. To assess whether the slope declines are caused by deregulation, I construct the following falsification test. First, I reassign 'pseudo' deregulation event years to one-time contractors. I do this by conditioning on the number of years between an establishment's regulation event and its last year in the data, and then randomly assign an 'age' for each establishment's pseudo event, taking draws from the conditional age distribution for the actual events. Then I reestimate the deregulation event study using these pseudo events. If the slope change in Panel B of Figure 1.3 is due to 'age' rather than deregulation *per se*, the same slope change should be evident in the pseudo event study. If slope change is due to deregulation, the slope change should be significantly less pronounced.

I plot the coefficients for the pseudo deregulation event study in Panel A of Figure 1.11. There is no discernible slope change following the pseudo event, which suggests that the slope change in Panel B of Figure 1.3 is indeed due to the deregulation.

Next, I test whether the degree of persistence depends on an establishment's experience as a contractor. I will return to this qualitative finding after introducing a model of endogenous screening capital below. I split contractors into two groups, those with 6 or fewer years between their regulation and deregulation events and those with more than 6 years between events, and estimate deregulation event studies for each group. Establishments in the first group have had an average of 2 years as contractors prior to their deregulation event, while establishments in the second group have had an average 8 years. In these event study models, I extend the endpoint b to 10 years following deregulation, and use the full deregulation sample for power. The results are plotted in Panel B of Figure 1.11. There are two things to note. First, the initial slope is higher for the experienced group (0.286 percentage points per year) than for the novice group (0.214 percentage points). Second, black share growth following deregulation is more persistent for the experienced group in both absolute and relative terms.

Finally, I assess the long run black share gains associated with AA regulation. I estimate a regulation event study with the endpoint b extended to 20 years and using the full regulation sample. Figure 1.12 displays the results. The point estimates are increasing up to 16 years after the regulation event, and then bounce around 2.4 percentage points, though the confidence intervals are relatively wide in this range. The black share gains associated with regulation remain quite apparent in the long run.

1.4 Causal Channels

A temporary AA program may produce a persistent increase in black employment through several causal channels. As described earlier, a temporary policy may increase black human capital investment, thereby reducing skill gaps (Coate and Loury 1993). Given the variation in regulation exploited here, it is unlikely the results here are driven by changes in human capital accumulation. The regulation of a particular firm should have little influence on incentives faced by workers, who

Figure 1.11: Fadeout Following Deregulation

(A) Pseudo Deregulation Event Study

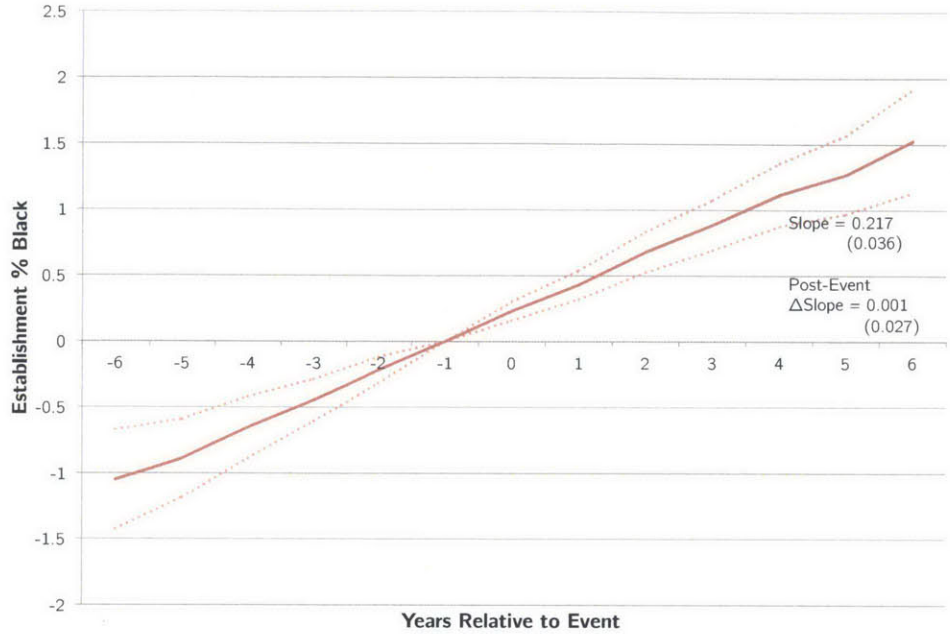

(B) Deregulation Event Study by Duration

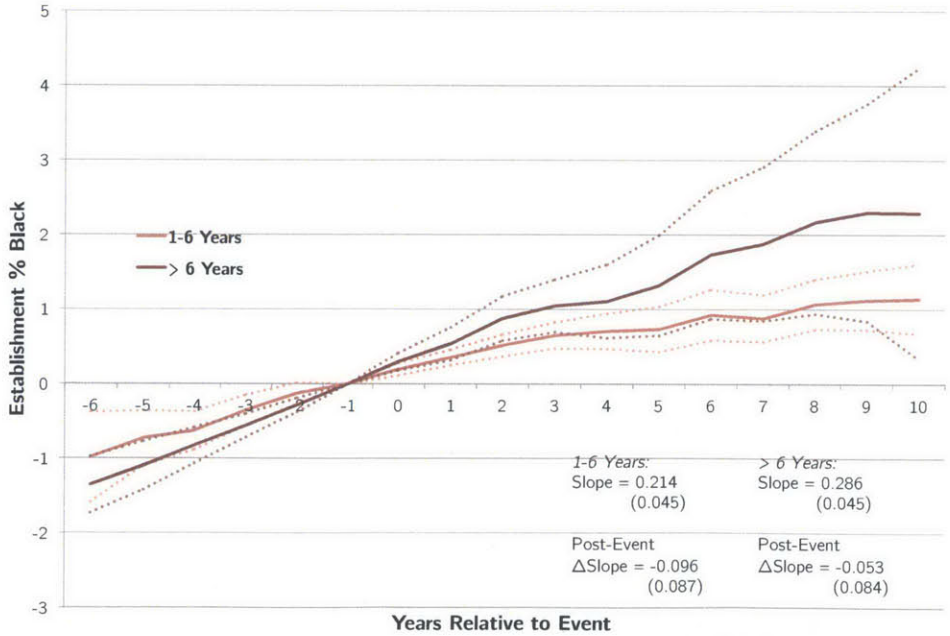

Notes: Notes: This figure plots event study coefficients and 95% confidence intervals (dotted) estimated using model 1.1 and the deregulation sample, where the outcome variable is the percent black of an establishment's employees. In Panel A, one-time contractors are assigned pseudo deregulation event years, as described in section 1.3.8. In Panel B, one-time contractors are grouped by the number of years between their regulation and deregulation events. The definitions of regulation and deregulation events are described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

Figure 1.12: Long Run Regulation Event Study

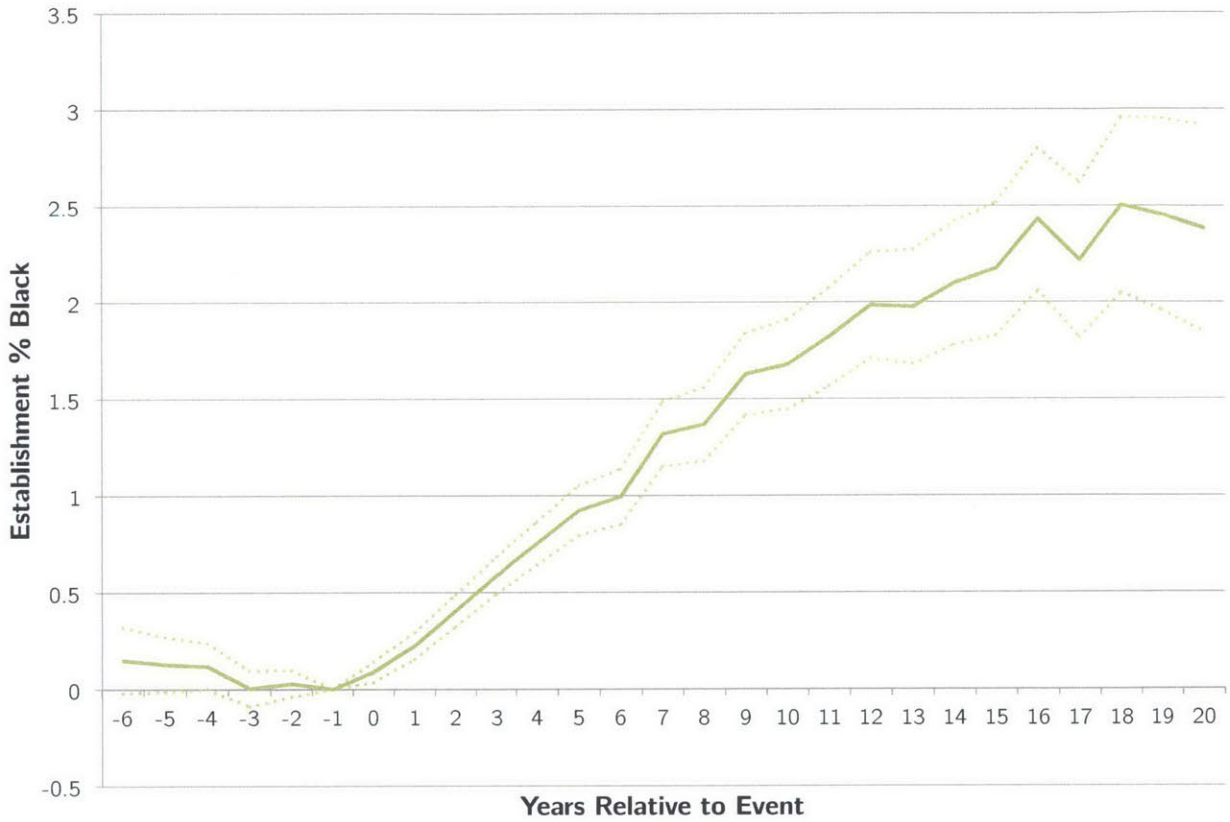

Notes: This figure plots regulation event study coefficients and 95% confidence intervals (dotted) estimated using model (1.1) and the regulation sample, where the outcome variable is the percent black of an establishment's employees and endpoint b is extended to 20. The definition of the regulation event is described in section 1.3.2. The coefficient for the year prior to the event (θ_{-1}) is normalized to zero. Estimated models include Census division by year fixed effects and a quadratic in log establishment size. Standard errors are clustered at the firm level.

presumably search in a broader labor market.

Temporary AA may spur employers to update negative stereotypes or reduce taste-based discrimination simply by increasing their exposure to black workers.[33] One may expect beliefs about worker quality to converge across employers. In that case, the impact of regulation on employer beliefs would dissipate over time as employers learned from other companies and updated their beliefs. Still, it is possible that these beliefs do not converge, or do so slowly, and that AA regulation continues to foster significant updating.

Temporary AA may also generate persistent black employment if employers face binding adjustment costs. Such costs may be relevant on a few margins. First, consider firing costs. If they are significant, a temporary policy may generate a long-lasting increase in the black share simply because workers hired while the policy was active are slow to leave. Given that I find that employee black share continues to *grow* after deregulation, firing costs alone cannot explain the results. However, firing costs in combination with some other mechanism could generate the persistence found here. For example, if the productivity of black potential workers is increasing in the black share of existing employees–say, due to mentoring (Athey et al. 2000) or role model effects (Chung 2000)–then temporary AA, in the presence of firing costs (or other adjustment costs), could produce the persistence found here.

Employers may also face adjustment costs when changing personnel practices. For example, a firm may formalize its hiring process as part of their AAP. They may maintain those formal practices following deregulation simply because it is costly to revert to an informal process. The lack of persistence found among police departments (Miller and Segal 2012) suggests that if adjustment costs are responsible for the persistence found here, they are somehow less significant or relevant for police departments.

For the remainder of the paper, I focus on the screening capital channel–AA may induce employers to make (partially) irreversible investments to improve screening. In this section, I formalize and discuss this channel in more detail.

1.4.1 A Screening Model with Endogenous Screening Capital

In this section I develop a simple screening model to explain the empirical results above. The model builds on the canonical Phelps (1972) model of statistical discrimination, and the setup is similar in spirit to Cornell and Welch (1996) and Morgan and Vardy (2009). In the model, an employer must hire a set of workers from a pool of candidates. The employer would prefer to hire the most productive candidates, but can only observe a noisy signal for each candidate's productivity. To improve its screening precision, the employer can invest in what I term *screening capital.* This term is meant to encompass all methods by which employers can improve screening. I interpret screening broadly as choosing the 'best' candidates from a set of potential workers, including both *recruitment*

[33] For example, Boisjoly et al. (2006) show that white college students who are randomly assigned black roommates are significantly more likely to endorse affirmative action and have personal contact with members of other ethnic groups years after assignment.

and *selection* components of the hiring process, which are often difficult to disentangle conceptually. Examples may include: employing and/or training personnel specialists and departments, developing job tests, developing relationships with and utilizing intermediaries such as employment agencies and schools, harnessing referral networks, or even learning by doing or experimentation. I show that if employers are initially less able to infer the quality of minority candidates–an assumption often made in the statistical discrimination literature–screening investments will decrease the hiring gap between groups. I discuss evidence justifying this assumption below.

I then introduce an AA regulation that constrains hiring rates to be equal across groups. I show that, under the same conditions, this regulation will *increase* the return to screening capital. The intuition is that screening investments generate more substantial improvements in the expected quality of minority hires, and affirmative action compels the employer to hire more minority workers.

Finally, if these screening capital investments are at least partially irreversible, then even a temporary AA regulation may generate persistent changes in hiring.

Suppose an employer must fill a mass n of vacancies from a mass $a(n) > 2n$ of applicants.[34] The wage is fixed.[35] There are two groups of candidates: $\gamma \in \{B, W\}$. Let π_γ denote the share of applicants from group γ. Worker match productivity is distributed

$$\theta \sim N\left(\mu_\theta(\gamma), \frac{1}{h_\theta}\right)$$

where the mean, $\mu_\theta(\gamma)$, is allowed to differ by group.

Neither the employer nor worker can observe a worker's match productivity directly. Instead, for each candidate i from group γ they observe a noisy signal for match productivity

$$s_i^\gamma = \theta_i + \epsilon_i^\gamma$$

where $\epsilon_i^\gamma \sim N\left(0, \frac{1}{h_\gamma}\right)$. Note that the signal precision, h_γ, is also allowed to vary by group.

Conditional on group γ and signal s, the expected productivity for a given applicant is

$$\mu(\gamma, s) = \frac{s h_\gamma}{h_\gamma + h_\theta} + \frac{\mu_\theta(\gamma) h_\theta}{h_\gamma + h_\theta}.$$

That is, expected productivity is a weighted average of the group mean, $\mu_\theta(\gamma)$, and the signal, s, where the weights depend on the precision of the signal relative to productivity dispersion.

Let $F(\hat{\mu}; \gamma)$ denote the cumulative distribution function for candidate's expected productivity

[34] All that is required here is that there are sufficient applicants so that the employer only hires from the right tail of the expected productivity distribution for each group.

[35] Suppose, for example, that the employer commits to a posted wage. This assumption follows both Cornell and Welch (1996) and Morgan and Vardy (2009).

after the signal realization. Then,

$$F(\hat{\mu};\gamma) = \Phi\left(\frac{\hat{\mu} - \mu_\theta(\gamma)}{\sqrt{\frac{1}{h_\theta}\frac{h_\gamma}{h_\theta + h_\gamma}}}\right)$$

and the expected productivity distribution for the entire applicant pool, $\bar{F}(\hat{\mu})$, is given by

$$\bar{F}(\hat{\mu}) = \pi_B F(\hat{\mu}; B) + \pi_W F(\hat{\mu}; W).$$

Given that the wage is fixed, the employer will simply hire the candidates with the highest expected productivity. Hence, it will set hiring threshold μ^* such that $1 - \bar{F}(\mu^*) = \frac{n}{a(n)}$.

Now, I will compare hiring rates for the two groups. Let $\lambda(\gamma) = 1 - F(\mu^*; \gamma)$ denote the hiring rate for group γ. Under an assumption often made in the statistical discrimination literature, we can sign the difference in hiring rates.

Proposition 1.4.1 *Suppose*

$$\mu_\theta(B) \leq \mu_\theta(W) \tag{1.6}$$

and

$$h_B < h_W \tag{1.7}$$

Then $\lambda(B) < \lambda(W)$.

In other words, if one group can be screened more precisely, that group will have an advantage at hiring.

The latter condition may hold if, for example, the primary screening method that the employer is endowed with is using referral networks, which tend to display group homophily (McPherson et al. 2001). Alternatively, workers may be better able to screen candidates from their own group (Giuliano et al 2009; Aslund et al. 2012).[36] In those cases, the employer may be endowed with a screening technology that favors the group better represented at the employer by some measure, for example, the group membership of the entrepreneur or owner.

Screening Capital

In the preceding section, the employer's ability to screen candidates was fixed. Suppose now that the employer can invest in screening capital to improve screening. In particular, suppose the employer can now pay cost $c(k)$ for an additional signal, s_i^k, where

$$s_i^k = \theta_i + \epsilon_i^k$$

where $\epsilon_i^k \sim N\left(0, \frac{1}{k}\right)$. Moreover, let $c'(k) > 0$ and $c''(k) > 0$. That is, the employer can pay cost $c(k)$ for signal with precision k, where the cost is increasing and convex in k. Note that the additional

[36] This may be due in part to differences in nonverbal listening and speaking cues, as in Lang (1986).

signal is equally informative about workers from both groups.[37]

Let $F(\hat{\mu};\gamma,k)$ denote the cumulative distribution function for candidates' expected productivity after the signal realizations for a level of screening capital k. Then,

$$F(\hat{\mu};\gamma,k) = \Phi\left(\frac{\hat{\mu}-\mu_\theta(\gamma)}{\sqrt{\frac{1}{h_\theta}\frac{h_\gamma+k}{h_\theta+h_\gamma+k}}}\right).$$

Under the conditions from Proposition 1.4.1 these investments will reduce hiring rate disparities.

Proposition 1.4.2 *Suppose (1.6) and (1.7) hold. Then*

$$\lambda(W,k)-\lambda(B,k)>0$$

is decreasing in k.

The intuition is that, with $h_B < h_W$, the additional signal will be more informative for B workers.

Now, consider the employer's problem when screening capital is available. Again, conditional on k, the employer will choose the candidates with highest expected productivity. Define $\mu^*(k)$ such that $1-\bar{F}(\mu^*(k);k)=\frac{n}{a(n)}$. Then, in choosing k, the employer's problem is

$$\max_k n\int_{\mu^*(k)}^{\infty}\mu\times\bar{f}(\mu;k)d\mu - c(k).$$

Denote the employer's solution by k^*.

AA Regulation

Suppose an AA regulation is introduced in this setting. I model the regulation as a mandate that the employer must equalize hiring rates across groups. Now the employer's optimal strategy is to choose two hiring thresholds: μ_B^* and μ_W^*. In particular, the employer will set $\mu_B^*(k)$ and $\mu_W^*(k)$ such that, for $\gamma\in\{B,W\}$,

$$1-F(\mu_\gamma^*(k);\gamma,k)=\frac{n}{a(n)}$$

In choosing k, the employer's problem is now

$$\max_k n\left[\pi_B\int_{\mu_B^*(k)}^{\infty}\mu f(\mu;B,k)d\mu+\pi_W\int_{\mu_W^*(k)}^{\infty}\mu f(\mu;W,k)d\mu\right]-c(k).$$

Proposition 1.4.3 *Suppose (1.6) and (1.7) hold. Then k^* is larger under AA.*

The intuition is that screening investments generate more substantial improvements in the expected quality of minority hires, and affirmative action compels the employer to hire more minority

[37] A richer model could allow the employer to choose among capital that provides signals more informative for one group than the other.

workers.

Finally, suppose these screening capital investments have sunk costs. Then in a dynamic setting, even a temporary AA regulation may generate persistent increase in screening capital, and hence produce a durable reduction in between-group hiring rate disparities.

1.5 Screening Evidence

The screening model outlined above generates two main testable predictions. First, AA will increase the return to screening capital. Second, screening investments will reduce between-group disparities in hiring rates.

Combining the EEO-1 data with additional establishment-level survey data, I provide evidence supporting the model. First, I show that regulated employers use more screening methods than otherwise comparable unregulated employers, particularly 'formal' methods such as personnel specialists, job tests, credential checks, and intermediaries. These results largely echo those in Holzer and Neumark (2000b).

Second, I show that an alternative source of variation in screening method use, *employer size*, also tends to equalize group representation among employees. Consistent with previous literature, I find that larger employers use more of the same screening methods. Moreover, I find that black share is increasing in employer size, while the *opposite* is true for black-run businesses.

Also note that a dynamic model would predict that employers expecting a longer spell of regulation would invest in more screening capital, all else equal. The intuition is that such an employer can spread (sunk) investment costs over more periods while the alternative means of compliance, lowering the hiring threshold for the covered group, entails an ongoing expense. This prediction further implies that employers with longer ex-ante regulation spells will display more persistent changes in hiring behavior following deregulation, consistent with findings discussed in 1.3.8.

In the next subsection, I introduce the data I use to study screening methods, the Multi-City Study of Urban Inequality (MCSUI).

1.5.1 Data

The Multi-City Study of Urban Inequality (MCSUI) is a cross-sectional survey of households and employers administered between June 1992 and May 1994 in four metropolitan areas: Atlanta, Boston, Detroit, and Los Angeles. I use the employer survey here, which includes over 3200 employers. The survey was conducted over the phone with individuals responsible for hiring at each establishment. The content focused on: establishment characteristics, including size and whether the establishment practices AA; employee characteristics, including the racial composition of employees and the last employee hired; and screening methods used, both generally and those used to select the most recent hire. The screening methods asked about include: whether the employer employs someone who specializes in hiring; whether the employer uses job tests, including skill tests, drug tests, and physicals; whether the employer verifies education credentials; and whether the em-

ployer checks criminal records; whether the employer uses intermediaries such as state, community, or private agencies, schools or unions to find job candidates; whether the employer uses employee referrals to find job candidates.

1.5.2 Affirmative Action and Screening Capital

As a *prima facie* test of the model, I first assess whether AA employers indeed invest more in screening capital than non-AA employers. I do this by testing whether AA employers use more screening methods. This exercise is in large part a replication of Holzer and Neumark (2000b).

To test this, I estimate models of the form

$$\text{method}_i = \alpha + \sum \lambda_{m(i)} + \beta AA_i + \gamma \log(\text{size})_i + X_i\delta + \epsilon_i \tag{1.8}$$

where i indexes employers, $\lambda_{m(i)}$ are MSA by central city fixed effects, and AA_i is an indicator for whether the employer reports accounting for affirmative action in recruiting or hiring.[38] I also include controls for other employer characteristic (namely, industry fixed effects), X_i, and the employer's number of employees, *size*. I also report the coefficient on size because an existing literature establishes a robust link between employer size and screening intensity, particularly the use of formal screening methods (Marsden 1994; Barron et al. 1987). I will further discuss and exploit this relationship below.

I present results for various screening methods in Table 1.6. AA employers are more likely to adopt a number of screening methods, and the differences are substantial. Conditional on covariates, AA employers are 7.1 percentage points more likely to use a personnel specialist, 4.8 percentage points more likely to use skill tests, 13.5 percentage points more likely to check education credentials, 7.8 percentage points more likely to check criminal backgrounds, and 11.9 percentage points more likely to use intermediaries. These differences are all significant at the 0.01 level.

In the final two columns I present estimates of (1.8) where the outcome is an indicator for whether the most recent hire was a referral. In column (6) I restrict estimation to establishments where the most recent hire was white, and in column (7) I restrict estimation to all other establishments.[39] When the last hire is white, that hire is 9.1 percentage points less likely to be referral at AA employers. When the last hire is non-white, that difference is less than 1 percentage point and statistically insignificant. Pooling all establishments, the last hire is 6.0 percentage points less likely to be a referral at AA employers. These estimates suggest that the use of these additional screening methods alters how hires are selected and therefore likely affects the composition of hires, *even within white hires*. The estimates also consistent with a particular interpretation of the screening model, where referrals are the endowed screening method; as employers invest in screening capital,

[38] Holzer and Neumark (2000b), the former author a principal investigator of the MCSUI survey, argue that this indicator primarily reflects variation in federal contractor status. The indicator may also include some employers with voluntary affirmative action plans.

[39] In contrast to the other results in this section, Holzer and Neumark (2000b) do not conduct an analysis similar to this.

Table 1.6: Affirmative Action and Screening Methods

	Personnel Specialist† (1)	Skill Test (2)	Check Education (3)	Check Crim. Record (4)	Use Intermediary (5)	Last Hire Referral White Hire (6)	Last Hire Referral Minority Hire (7)
Affirmative Action	0.071**	0.048**	0.135**	0.078**	0.119**	-0.091**	-0.009
	(0.016)	(0.017)	(0.020)	(0.019)	(0.019)	(0.028)	(0.029)
log(Size)	0.121**	0.011*	0.035**	0.038**	0.043**	-0.034**	-0.041**
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
1-Digit Industry FEs	✓	✓	✓	✓	✓	✓	✓
MSA by CC FEs	✓	✓	✓	✓	✓	✓	✓
N Firms	2766	2757	2546	2514	2766	1418	1302
R^2	0.300	0.038	0.113	0.142	0.080	0.049	0.051
Mean of Dep. Var.	0.325	0.239	0.377	0.323	0.634	0.375	0.353

Notes: Each column reports the coefficient estimates for a regression of the form (1.8), with robust standard errors in parentheses. Data are from the Multi-City Study of Urban Inequality (MCSUI). Affirmative Action status is defined by whether the establishment reports that "Affirmative Action or Equal Opportunity LawÕÕ plays any role in recruiting or hiring. For each column, the outcome variable is a particular screening method. 'Personnel Specialist' is an indicator for whether the establishmentÕs most recent search was carried out by a personnel official. The outcome variables for columns (2)-(5) are indicators for whether the indicated method is typically used in screening. 'Last Hire Referral' is an indicator for whether the last employee hired was an employee referral. Column (6) restricts estimation to establishments where the last hire was white, while column (7) restricts estimation to all other establishments. All models include 1-digit industry fixed effects and MSA by central city status fixed effects. Establishments are weighted using survey sample weights.

* Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.

they become less reliant on referrals.[40]

1.5.3 Employer Size and Black Share

One additional point to note from Table 1.6 is that the coefficient on log size is the same sign as the coefficient on AA and statistically significant for each dependent screening variable. This is consistent with existing evidence that larger employers screen more intensely, use more screening methods, and rely less on referrals (Barron et al. 1987; Holzer 1987; Marsden 1994). In this next section, I exploit employer size as an additional source of variation in screening methods to test a prediction of the model.

If regulated employers increase their black share of employees in part by investing in screening capital, then other exogenous sources of variation in screening capital across employers should also predict variation in black share in the same manner. I show above that, consistent with previous research, larger employers utilize more screening methods. As another test of the model, I examine whether employer size predicts an employer's black share.

There are three types of explanations for the observed relationship between employer size and

[40]This result does *not* imply that employers do not use referrals to comply with AA regulation. In particular, employers may change the *composition* of referrals they hire. Referrals may also play a key reinforcing role in compliance and persistence, even if they are not the employer's initial focus. Any initial push to employ more minority workers may diversify the pool of potential referral hires, allowing the employer to better screen minority candidates in future hiring.

screening intensity. First, many screening methods may be cheaper for larger employers per-hire basis due to economics of scale. For example, employing a personnel specialist or developing a skill test may involve substantial fixed costs. Second, larger employers may face a higher return to worker quality and so put more emphasis on screening (Barron et al. 1987).[41] Third, larger employers may face additional pressure from the EEOC; indeed, Title VII of the Civil Rights Act of 1964, which outlaws workplace discrimination, only applies to firms with at least 15 workers. Larger employers may adopt formal screening procedures to be in compliance, either by actually employing a more diverse workforce or by adopting personnel practices that are less susceptible to discrimination claims.

The first two explanations suggest that variation in screening capital deriving from employer size may be taken as exogenous. However, the third explanation is problematic because it suggests a confounding relationship between employer size and black share. In particular, larger firms may face pressure to both employ more minority workers and to use formal screening methods, even if the latter does not lead to the former. I will come back to this issue later and present additional evidence that suggests that EEO pressure is unlikely to be the primary driver of the relationship between employer size and black share that I observe.

To estimate the relationship between employer size and black share, I first use the EEO-1 data. I estimate models of the form

$$(\text{black share})_{it} = \alpha_i + \lambda_{d(i),t} + \beta^{est} \log(\text{establishment size})_{it} + \beta^{firm} \log(\text{firm size})_{it} + \epsilon_{it}. \quad (1.9)$$

I include firm size as an explanatory variable in some specifications because whatever causes larger establishments to do more worker screening may produce a similar relationship between *firm* size and screening. Hence, black share may also be increasing in firm size. I measure firm size as the total number of employees at establishments reported in the EEO-1 data under the same firm. Note that black share is still measured at the establishment level.

Results are presented in panel A of Table 1.7. All models include Census division by year fixed effects. Columns (2) and (3) include establishment fixed effects. Columns (5) and (6) use only within-job variation, where jobs are defined by establishment by occupation cells. In all models, establishment and firm size are significant predictors of establishment black share. Surprisingly, including establishment fixed effects *increases* the coefficient on log establishment size. Isolating within-job variation reduces the coefficient on log establishment size by more than half, but the coefficient remains sizable. Larger employers do employ more workers in occupations that black workers tend to work in, but this alone cannot explain the relationship. The coefficients from column (5) imply that a 10% increase in establishment size predicts a 0.07% percentage point increase in the black share of employees within jobs, while a 10% increase in *firm* size predicts a 0.02% percentage point increase in the black share within jobs. Hence, the relationship is primarily

[41] Indeed, this may be why large employers are large in the first place. Alternatively, larger employers may face higher monitoring costs (Barron et al. 1987). This may also increase the return to worker quality, defined appropriately (e.g. work ethic).

Table 1.7: Employer Size and Black Share

Outcome: Percentage Black				Within-Job	
	(1)	(2)	(3)	(4)	(5)
log Establishment Size	1.434	1.659	1.617	0.666	0.568
	(0.027)	(0.035)	(0.037)	(0.025)	(0.027)
log Firm Size			0.104		0.227
			(0.028)		(0.025)
Establishment FEs		✓	✓	✓	✓
MSA by Year FEs	✓	✓	✓	✓	✓
R^2	0.230	0.927	0.927	0.872	0.872
Outcome: Δ Percentage Black			Within-Job		
	All	Recession	Expansion	Non-Contractors	Contractors
	(1)	(2)	(3)	(4)	(5)
Δ log Establishment Size	0.781	0.746	0.801	0.873	0.757
	(0.024)	(0.037)	(0.030)	(0.053)	(0.027)
MSA by Year FEs	✓	✓	✓	✓	✓
R^2	0.002	0.001	0.002	0.002	0.002

Notes: Each column reports the coefficient estimates for a regression, with standard errors in parentheses clustered at the establishment level. In Panel A the outcome variable is the percent black of an establishmentÕs employees (or, as in columns (5) and (6), the percent black of an establishment by occupation cell). Firm size is the total number of employees at establishments reported in the EEO-1 data under the same firm. Columns (2) and (3) include establishment fixed effects. Columns (4) and (5) isolate within-job variation in black share. In Panel B the outcome variable is the change in percent black of employees in a establishment by occupation cell over the previous year. All columns include MSA by year fixed effects. Column (2) includes only data from the years 1980-1982, 1990-1992, and 2001-2003. Column (3) includes only the remaining years. Column (4) includes only observations for establishments that have not previously held a federal contract. Column (5) includes only observations for establishments that have previously held a federal contract.

establishment level.

Next, I explore the robustness of this result. I test several alternative explanations for the positive relationship between employer size and black share found here. First, I test whether this relationship is an artifact of the business cycle. For example, if establishments tend to grow during expansions, and black job seekers make up a larger fraction of the applicant pool during expansions, then employers will tend to increase their black share as they grow. Second, I test whether this relationship is due to AA. I find above that AA causes the black share of employees to increase. If this is primarily driven by regulated employers increasing their black share of new hires, for example, then the black share may increase more for growing establishments. Third, I test whether this relationship is generated by changes in the composition of occupations. For example, larger establishments may tend to require more workers in low skill occupations, while black workers make up a disproportionately large share of workers in those occupations.

I test these alternative hypotheses, focusing on within-job changes in black share. I estimate models of the form

$$\Delta(\text{black share})_{iot} = \alpha + \lambda_{d(i),t} + \beta\Delta\log(\text{establishment size})_{it} + \epsilon_{iot} \tag{1.10}$$

where $\Delta x_{it} = x_{i,t} - x_{i,t-1}$. I estimate this model using all the data, separately for recession and expansion years, and separately for contractors or one-time contractors and establishments with no contractor experience. Note that I measure size changes at the *establishment* level, not the job cell level.

The results are presented in panel B of Table 1.7. Using the full data, the first difference model produces a β coefficient of 0.781-that is, a 10% increase in establishment size predicts about a 0.08 percentage point increase in the black share of employees within jobs. Estimates are comparable during both economic recessions and expansions, contradicting the business cycle hypothesis. Finally, the relationship between employer size and black share is even *larger* for establishments with no contractor experience. This is consistent with AA regulation inducing employers to make screening investments they would otherwise make as larger establishments.

Though a single β coefficient can summarize the relationship between job black share and size, this relationship may vary with size. To explore this relationship in more depth, I estimate models of the form

$$(\text{black share})_{iot} = \alpha_{io} + \lambda_{d(i),t} + \sum_{s=1}^{S} \beta_s I_{it}^s + \epsilon_{iot} \tag{1.11}$$

where I^s are indicators for the following establishment size bins: 5-50, 51-100, 101-150, ..., 451-500, and 501+. Critically, the EEO-1 data allow me to measure the relationship between employer size and job black share using *within-establishment* variation in size. This is in contrast to a small existing literature documenting a positive cross-sectional relationship between establishment size and black share (Holzer 1998; Carrington et al. 2000). In this literature, it is unclear whether this cross-sectional relationship is driven by fixed differences between employers. For example, one explanation posited in Holzer (1998) is that larger firms tend to be located in more urban locations, where minorities tend to work. Carrington et al. (2000) use firm size to identify the causal effect of Title VII. By including establishment fixed effects, I can remove this type of variation. Moreover, all establishments in these data are covered by workplace discrimination laws. Hence EEO law coverage alone cannot explain any relationship found here.

I plot the β_s coefficients in Figure 1.13. I plot two sets of coefficients, one set estimated using only establishments that have previously held a contract, and the other using only establishments that have not. I also include 95% confidence intervals for both sets of coefficients. For both sets, there is a clear positive relationship between establishment size and minority size. Relative to establishments with 5-50 employees, growing to 200-250 employees is associated with a black share increase of nearly 1.5 percentage points. Growing further to 500 or more employees is associated with an additional increase of 0.5 percentage points.

Black-Run Businesses

The model also predicts that screening investments will have different effects depending on the initial composition of the employer. For example, if an employer's hiring manager is endowed with a higher ability to screen black job candidates than white job candidates, this would reverse the predictions

Figure 1.13: Establishment Size and Job Black Share

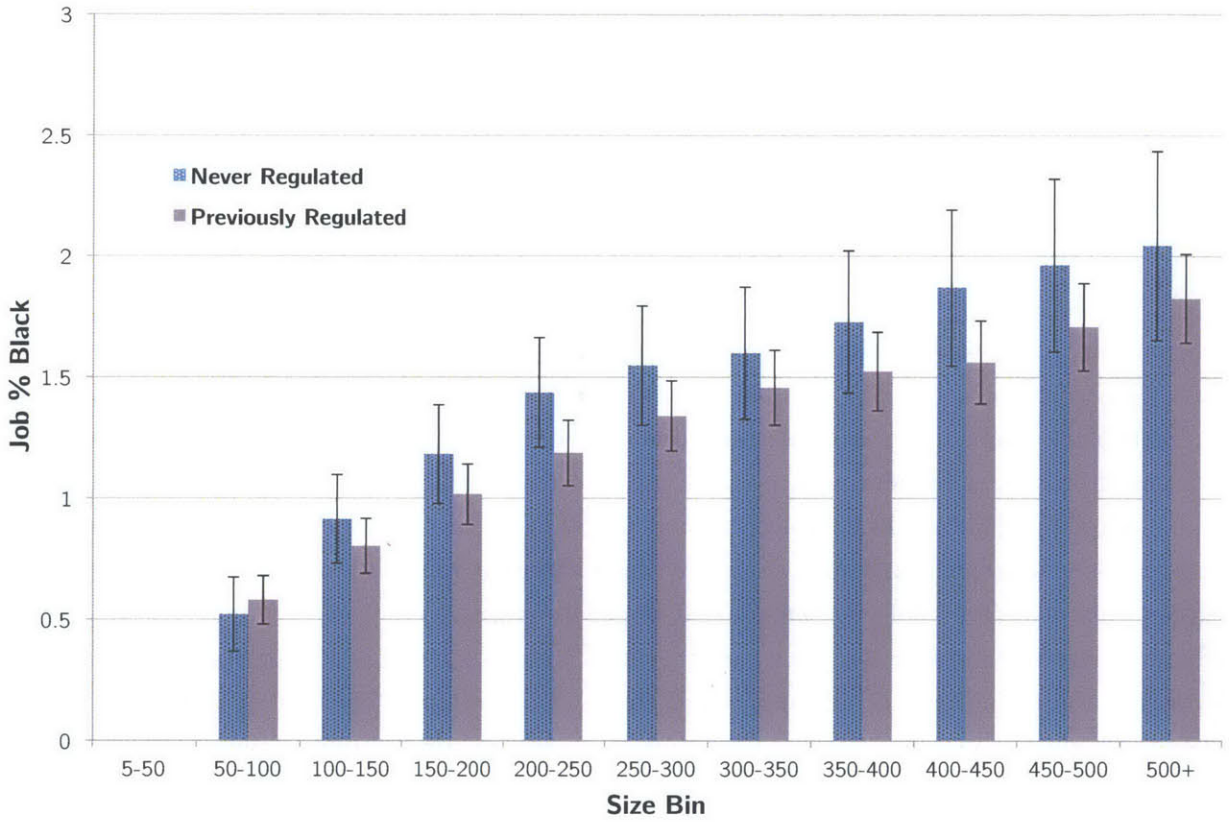

Notes: This figure plots coefficients and 95% confidence intervals from regression where the outcome variable is the percent black of employees in an establishment by occupation cell and the explanatory variables of interest are a set of indicators for establishment size bins. All models include Census division by year fixed effects. The purple bars depict coefficient estimates from a model that includes only establishments that have previously held a federal contract. The blue bars depicts estimates from a model that includes only establishments that have not previously held a contract. Standard errors are clustered at the establishment level.

of the model: an investment in screening capital would reduce the gap in screening precision between groups, increasing the rate at which white candidates are hired. In general, screening investments should reduce between-group differences in hiring rates. By contrast, EEO law is historically unlikely to pressure private employers to hire white workers. Hence, if the relationship between employer size and black share were entirely driven by EEO pressure, the relationship should not change sign depending on the employer's endowed screening technology.

To test this prediction, I assess whether the size-minority share relationship is of opposite sign for minority-run businesses. The motivating intuition is that minority-run businesses should be better endowed to screen minority candidates than white-run businesses. However, as these employers grow they invest in screening capital and their endowed screening ability becomes less relevant. In particular, using the MCSUI data I estimate models of the form

$$(\text{black share})_i = \sum_{j \in \{B,W\}} \alpha^j + \lambda^j_{m(i)} + \beta^j \times \log \text{size} X_i \gamma^j + \epsilon_i \tag{1.12}$$

where j indexes whether the establishment's hiring manager is white or black.

The results are presented in Table 1.8. Column 1 does not include industry controls, column 2 includes 1-digit industry fixed effects, and column 3 includes 2-digit industry fixed effects. The results are similar across columns. As predicted, while establishment black share is increasing in establishment size for white-run businesses, the *opposite* is true for black-run businesses. For white-run and black-run establishments, a 10 log point increase in size is associated with a 0.25-0.26 percentage point increase and a 0.40-0.54 percentage point decrease in black share.

Black Share and Discrimination Charges

If larger employers face more pressure to comply with workplace discrimination law (or if they are simply more aware of compliance issues), they may hire more minority workers to satisfy law enforcement officials and reduce costs associated with discrimination claims and litigation. In this case, workplace discrimination law may be responsible for the positive relationship between employer size and black share. However, it is not clear that hiring more minority workers will reduce discrimination claims and associated costs, particularly given that claims are typically filed by individuals. Only a small fraction discrimination claims relate to hiring, especially after the mid-1970's (Donahue and Siegelman 1991). Hence, nearly all claims are filed by individuals who were employees at some point.

I estimate the relationship between an establishment's black share and the expected number of discrimination charges this establishment faces in a given year. To do this, I use data on the universe of discrimination claims filed with the Equal Employment Opportunity Commission from 1990 to 2004. These data include information for each charge filed, including the protected class of the complainant, the nature of the allegation, how that claim is resolved, and identifying information for the accused employer. I restrict analysis to claims alleging racial discrimination against black workers filed under Title VII of the Civil Rights Act of 1964. I match these claims to the EEO-1

Table 1.8: Manager Race, Size, and Black Share

	Percent Black		
	(1)	(2)	(3)
log Size × White-Run	2.50**	2.54**	2.63**
	(0.485)	(0.512)	(0.531)
log Size × Black-Run	-5.39**	-5.42**	-4.11*
	(1.371)	(1.510)	(1.805)
1-Digit Industry by Race FEs		✓	
2-Digit Industry by Race FEs			✓
MSA by CC by Race FEs	✓	✓	✓
N White-Run Establishments	2166	2124	2124
N Black-Run Establishments	198	190	190
R^2	0.195	0.199	0.253

Notes: Each column reports the coefficient estimates for a regression, with robust standard errors in parentheses. Data are from the Multi-City Study of Urban Inequality (MCSUI). 'White-Run' and 'Black-Run' are indicators for whether the employee that oversaw the most recent search identifies as white or black. All models include MSA by central city status by race fixed effects. Column (2) includes 1-digit industry by race fixed effects. Column (3) includes 2-digit industry by race fixed effects.
* Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.

data using the employer name and address. I am able to match 140,124 claims. I present summary statistics in Table 1.C.2 in the Appendix.

I estimate Poisson models of the form

$$\log(E[\text{charges}_{it}]) = \alpha_i + f(\log(\text{size})_{it})\theta + g(\text{black share}_{it})\beta + \epsilon_{it} \tag{1.13}$$

where $f(\cdot)$ is a 10-piece linear spline, and g is a 5th order fractional polynomial. I estimate models that do and do not include establishment fixed effects, α_i. In addition to discrimination charges, I also estimate models where the dependent variable is the number of claims that result in what the EEOC terms 'merit resolutions'–charges with outcomes favorable to the claimant or where the EEOC finds that allegations have merit. These are charges that result in a withdrawal with benefits, settlement with benefits, successful conciliation, or unsuccessful conciliation. Of the claims I am able to match to the EEO-1 data, 14% are merit claims.

The estimated polynomials in black share are plotted in Figures 1.14. The black dashed lines mark the 10th and 90th percentiles of establishment black share in the data. Conditional on size, the expected number of racial discrimination charges against an establishment is *increasing* in black share at all but the highest black share levels. Hence, for nearly all establishments, hiring more black workers does not appear to reduce dealings with discrimination law enforcement. Moreover, the expected number of claims that result in merit resolutions is also *increasing* in black share at all but the highest levels.

Figure 1.14: Black Share and Discrimination Charges

(A) Discrimination Charges

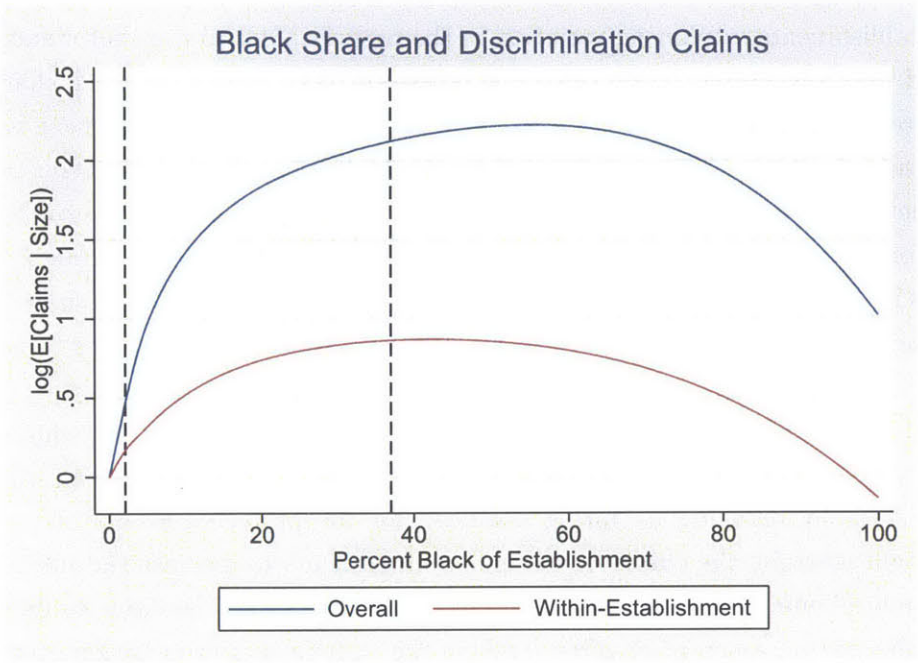

(B) Merit Resolutions

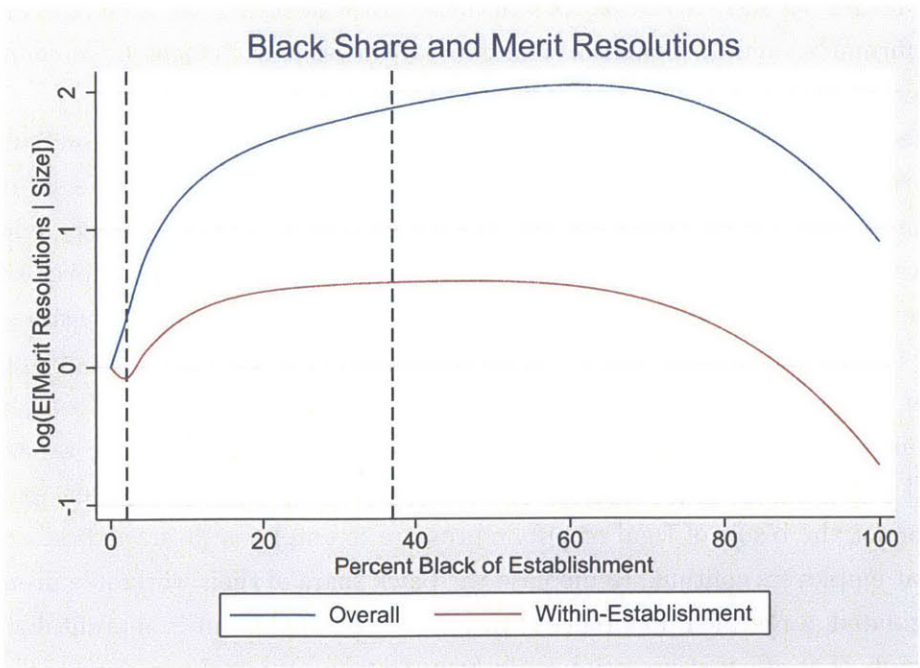

Notes: These figures plot the expected number of racial discrimination charges received by an establishment as a function of its black share of employees, conditional on establishment size. This relationship is estimated using Poisson regression models. The regressions include a 10-piece linear spline in log establishment size and a 5th order fractional polynomial in establishment black share. In both panels, the red line depicts a model that includes establishment fixed effects. The black dashed lines mark the 10th and 90th percentile establishment black shares in the data. In Panel A the dependent variable is the number of charges that include an allegation of discrimination against black workers. In Panel B the dependent variable is the number of merit resolutions that include an allegation of discrimination against black workers. Merit resolutions are discrimination charges that result in a withdrawal with benefits, settlement with benefits, successful conciliation, or unsuccessful conciliation.

1.6 Conclusion

I estimate the dynamic effects of federal affirmative action regulation on the racial composition of regulated establishments, exploiting variation in the timing of federal contractor status across workplaces. I find that affirmative action sharply increases an establishment's black share of employees, with the share continuing to increase over time. Five years after an establishment is first regulated, its black share of employees increased by an average of 0.8 percentage points. This response is strikingly asymmetric: even after establishments are deregulated, their black share *continues to grow*, in some cases for more than a decade. By contrast, establishment size moves roughly symmetrically with changes in 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.

I argue that this persistence is in part driven by affirmative action inducing employers to improve screening. In existing models (e.g., Coate and Loury 1993), an employer can only comply with affirmative action by reducing its hiring standard for the protected group. I introduce a novel response margin whereby the employer can make investments to increase the precision with which it screens potential employees. I show that if employers are initially less able to infer the quality of minority candidates–an assumption often made 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 predictions of the model. Using cross-sectional survey data, I show that regulated employers use more screening methods than otherwise similar unregulated employers, particularly formal methods like personnel specialists, job tests, credential checks, and intermediaries. Exploiting variation in screening capital deriving from employer size, I also show that screening capital predicts reductions in between-group employment disparities. In particular, I show that black share is increasing in employer size, and that this relationship holds *within-establishment* and *within-job* for a large sample of establishments that are all subject to workplace discrimination law. In this sense, affirmative action pushes employers to hire as though they were expanding. I provide several pieces of evidence that the relationship between employer size and black share is not the result of legal or public pressure to employ a diverse workforce.

Given that employers continue to increase the black share of their workforce 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, ideally using worker level data. Relatedly, these two classes of explanations have different implications for the productive efficiency consequences of affirmative action.

Note that, even in the absence of adjustment costs, the results presented here do not necessarily imply the existence of any market inefficiency in personnel practices. Screening capital may be costly, and employers may not find additional capital to be a worthwhile investment in the absence of regulation. If regulation prompts an employer to invest, in the presence of sunk costs, the employer may find the screening capital profitable to use even following deregulation. Nevertheless, the results do suggest that the distortionary costs of affirmative action may be smaller than previously assumed. In particular, screening capital may involve upfront costs, but little ongoing distortion in who is employed. Given that the persistence following deregulation appears to be incomplete, compliance may involve a combination of screening investment and a lowered quality threshold for black hires. A lowered threshold is an ongoing distortion, hence only the capital response would be persistent absent adjustment costs.

Another remaining question is how the compliance activities of one employer affect the aggregate labor market. Consider screening investments as one response. If more precise screening improves the matching of workers to jobs, then screening investments may improve productive efficiency. But if improved screening primarily redistributes workers among competing employers without increasing match quality, improvements in productive efficiency will be negligible (Stiglitz 1975). Moreover, screening improvements can potentially amplify adverse selection in the labor market (Masters 2009). Relatedly, screening improvements at one employer may increase compliance costs for other employers. Other compliance activities may exhibit similar crowd out or reallocation across employers (Donahue and Heckman 1991). Newly available employee-employer matched data provide a unique opportunity to identify the extent of this reallocation, a potentially useful input for determining the aggregate impact of affirmative action regulation.

In contrast to persistence produced through endogenous human capital investment by workers as in Coate and Loury (1993), the persistent effect of temporary affirmative action found here appears to be employer level rather than market level. A potentially important implication is that employer churn may lead the effect of temporary affirmative action to fade out in the aggregate. If the persistence found here is in part driven by the human capital accumulation of *employers*–say, from hiring managers learning how to screen minority workers–then this fadeout may be partially muted by the transfer of knowledge across employers.

The results presented here also have implications outside of affirmative action policy. More generally, they suggest 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.

1.A Appendix: Persistence Can Introduce Bias

The persistent effect of temporary regulation I document here has important implications for interpreting existing research in this literature, including Kurtulus (2011, 2012), Leonard (1984, 1990), Rodgers and Spriggs (1996), Ashenfelter and Heckman (1976), Goldstein and Smith (1976), Smith and Welch (1984), and Heckman and Wolpin (1976). In particular, if regulation has an impact on employers that persists even when they are no longer contractors, previous estimates may be biased. This is because the research designs applied in existing work are based on simple comparison of contractors to non-contractors, either within or across employers. In the presence of persistence these comparisons may substantially understate the causal impact of regulation because some employers that are currently non-contractors were previously contractors, and the minority share of those employers is still affected by the regulation. In this section, I explore the extent of this bias empirically, using the baseline model of Kurtulus (2011) as a motivating example.

The core specification estimated in Kurtulus (2011) is of the form

$$Y_{it} = \alpha_i + \tau_{r(i),t} + \beta I_{it}^{current} + X_{it}\gamma + \epsilon_{it} \tag{1.A.1}$$

where $I^{current}$ is an indicator for whether an establishment is currently a contractor. This specification models the effect of regulation as a level effect that depends only on the current period contractor status. An assumption implicit in this model is that whatever effect regulation has dissipates completely when an employer is no longer a contractor.

For the sake of comparison, I also estimate a modified version of (1.A.1),

$$Y_{it} = \alpha_i + \tau_{r(i),t} + \tilde{\beta} I_{it}^{previous} + X_{it}\gamma + \epsilon_{it} \tag{1.A.2}$$

where $I^{previous}$ is an indicator for whether an establishment has *ever* previously been a contractor. This specification models the effect of regulation as a level effect that depends only on whether the establishment was ever previously a contractor. While this specification does not allow the effect of regulation to accumulate over time, a pattern I document in the main analysis, it does allow for a particular form of persistence. If the effect of regulation takes the form assumed in (1.A.1), this model will underestimate the effect of regulation.

As in the main analysis, I exclude establishments that enter the sample as a federal contractor. In addition, for the establishments that become contractors, I only include years of data that are at most 6 years prior to their regulation event. To demonstrate the influence of persistence on the results, I estimate both (1.A.1) and (1.A.2) for a series a estimation samples, moving the data window from the year of the regulation event to 6 years following the regulation event. For each sample, I also restrict the set of eventual contractors included in the estimation to those that are present for the full set of years following the regulation event. This way, the $\tilde{\beta}$ estimates reflect, in principle, the impact of regulation averaged across 0-6 years following the event, and not a more complicated weighted average that depends on the frequency with which establishments are observed

at each year following regulation. Each model includes Census division by year fixed effects, and a quadratic in log establishment size.

The β and $\tilde{\beta}$ estimates for each window are displayed in Figure 1.C.4. For the smallest window, which excludes all years following the regulation event, the estimates coincide at 0.141. As the window widens, these estimated coefficients diverge sharply. The β coefficient declines to 0.068 when the window expands to three years after the event, and is statistically indistinguishable from zero. Using the full size year window, the β coefficient declines further to 0.051. This pattern emerges despite the fact that the effect of initial regulation *increases* over time, as demonstrated in the main analysis. This discrepancy reflects the fact that many establishments are not contractors in some years following their regulation event, but their black share continues to increase. By contrast, the $\tilde{\beta}$ coefficient increases substantially as the window expands. With a three year post-event window, the coefficient has more than doubled to 0.340. Using the full size year window, the $\tilde{\beta}$ coefficient increases further to 0.388. A simple adjustment allowing for some form of persistence increases the estimated effect of regulation by an order of magnitude.

1.B Appendix: Proofs of Propositions

In this section I prove the propositions stated in section 1.4.1.

I begin by introducing some notation that will be helpful throughout. Define $z^*(\gamma)$ and $z^*(\gamma, k)$ as

$$z^*(\gamma) \equiv \frac{\mu^* - \mu_\theta(\gamma)}{\sqrt{\frac{1}{h_\theta}\frac{h_\gamma}{h_\theta + h_\gamma}}} \quad \text{and} \quad z^*(\gamma, k) \equiv \frac{\mu^* - \mu_\theta(\gamma)}{\sqrt{\frac{1}{h_\theta}\frac{h_\gamma + k}{h_\theta + h_\gamma + k}}}.$$

Proposition 4.1 *Suppose*

$$\mu_\theta(B) \leq \mu_\theta(W)$$

and

$$h_B < h_W$$

Then $\lambda(B) < \lambda(W)$.

Proof We have that

$$\begin{aligned}\lambda(\gamma) &= 1 - \Phi(z^*(\gamma)) \\ &= 1 - \Phi\left(\frac{\mu^* - \mu_\theta(\gamma)}{\sqrt{\frac{1}{h_\theta}\frac{h_\gamma}{h_\theta + h_\gamma}}}\right)\end{aligned}$$

By assumption, we have that $\mu^* \geq \mu_\theta(W) \geq \mu_\theta(B)$ and $0 < h_B < h_W$. Hence, $z^*(B) > z^*(W)$ and $\lambda(B) < \lambda(W)$.

Proposition 4.2 *Suppose (1.6) and (1.7) hold. Then*

$$\lambda(W,k) - \lambda(B,k) > 0$$

is decreasing in k.

Proof[42] Given that the total hiring rate is constant, a change in k must either leave group hiring rates unchanged or move them in opposite directions. Further,

$$\begin{aligned}
\frac{\partial \lambda(\gamma)}{\partial k} &= \frac{\partial}{\partial k}[1 - \Phi(z(\gamma,k))] \\
&= \phi(z(\gamma,k)) \left[\frac{1}{2} \frac{h_\theta}{(k+h_\gamma)(k+h_\gamma+h_\theta)} z(\gamma,k) - \frac{\partial \mu^* / \partial k}{\sqrt{\frac{1}{h_\theta} \frac{h_\gamma+k}{h_\theta+h_\gamma+k}}} \right] \\
&= \frac{\phi(z(\gamma,k))}{\sqrt{\frac{1}{h_\theta} \frac{h_\gamma+k}{h_\theta+h_\gamma+k}}} \left[\frac{1}{2} \frac{h_\theta}{(k+h_\gamma)(k+h_\gamma+h_\theta)} (\mu^* - \mu_\theta(\gamma)) - \frac{\partial \mu^*}{\partial k} \right].
\end{aligned}$$

Given that $f(B) > f(W)$, it can't be that $\frac{\partial \lambda(B)}{\partial k} = \frac{\partial \lambda(W)}{\partial k} = 0$. Hence, $\frac{\partial \lambda(B)}{\partial k}$ and $\frac{\partial \lambda(W)}{\partial k}$ must be opposite signed, and so $f(B) > f(W)$ implies $\frac{\partial \lambda(B)}{\partial k} > 0$ and $\frac{\partial \lambda(W)}{\partial k} < 0$.

Proposition 4.3 *Suppose (1.6) and (1.7) hold. Then k^* is larger under affirmative action.*

This claim can be restated as

$$\frac{\partial}{\partial k} \Delta > 0.$$

where

$$\Delta = \pi_B \int_{\mu_B^*(k)}^{\infty} \mu dF(\mu|B,k) + \pi_W \int_{\mu_W^*(k)}^{\infty} \mu dF(\mu|W,k) - \pi_B \int_{\mu^*(k)}^{\infty} \mu dF(\mu|B,k) - \pi_W \int_{\mu^*(k)}^{\infty} \mu dF(\mu|W,k).$$

To show this, I first change the variable of integration, expressing the integrals in terms of quantile functions rather than cumulative distribution functions. In particular, let $Q(p|\gamma,k) = \{\mu : p = F(\mu|\gamma,k)\}$. Then

$$\begin{aligned}
\pi_\gamma \int_{\mu_\gamma^*(k)}^{\infty} \mu dF(\mu|\gamma,k) &= \pi_\gamma \int_{F(\mu_\gamma^*|\gamma,k)}^{1} Q(p|\gamma,k) dp \\
&= \pi_\gamma \int_{\alpha}^{1} Q(p|\gamma,k) dp
\end{aligned}$$

[42]The proof is similar to Claim III.D.1 in Autor and Scarborough (2008).

where $\alpha = \frac{n}{a(n)}$. In these terms, we can express Δ as

$$\begin{aligned}\Delta &= \pi_B \int_{\alpha}^{1} Q(p|B,k)dp + \pi_W \int_{\alpha}^{1} Q(p|W,k)dp - \pi_B \int_{\alpha_B}^{1} Q(p|B,k)dp - \pi_W \int_{\alpha_W}^{1} Q(p|W,k)dp \\ &= \pi_B \int_{\alpha}^{\alpha_B} Q(p|B,k)dp - \pi_W \int_{\alpha_W}^{\alpha} Q(p|W,k)dp.\end{aligned}$$

Note that

$$Q(p|\gamma,k) = \mu_\theta(\gamma) + \sqrt{\frac{1}{h_\theta}\frac{h_\gamma + k}{h_\theta + h_\gamma + k}}\Phi^{-1}(p)$$

and so

$$\frac{\partial Q(p|\gamma,k)}{\partial k} = \sqrt{\frac{1}{h_\theta}}\Phi^{-1}(p)\frac{\partial}{\partial k}\left(\frac{h_\gamma + k}{h_\theta + h_\gamma + k}\right) > 0$$

which is decreasing in h_γ.

Hence,

$$\begin{aligned}\frac{\partial}{\partial k}\Delta &= \frac{\partial}{\partial k}\left[\pi_B \int_{\alpha}^{\alpha_B} Q(p|B,k)dp - \pi_W \int_{\alpha_W}^{\alpha} Q(p|W,k)dp\right] \\ &> \frac{\partial}{\partial k}\left[\pi_B \int_{\alpha}^{\alpha_B} Q(p|B,k)dp - \pi_W \int_{\alpha_W}^{\alpha} Q(p|B,k)dp\right] \\ &> 0\end{aligned}$$

where the first inequality holds because $\frac{\partial Q(p|\gamma,k)}{\partial k}$ is decreasing in h_γ and the second inequality holds because $\frac{\partial Q(p|\gamma,k)}{\partial k}$ is increasing in p.

1.C Appendix: Additional Figures and Tables

Figure 1.C.1: Summary Statistics by Event Study Year

(A) Number of Establishments in Event Studies

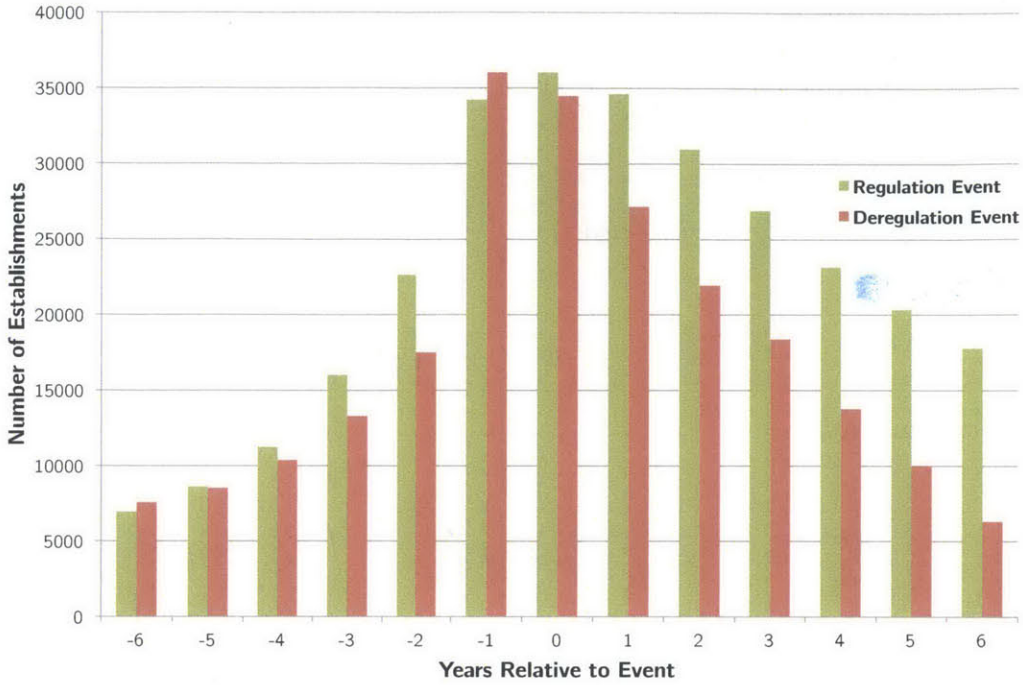

(B) Share Contractor by Year in Event Studies

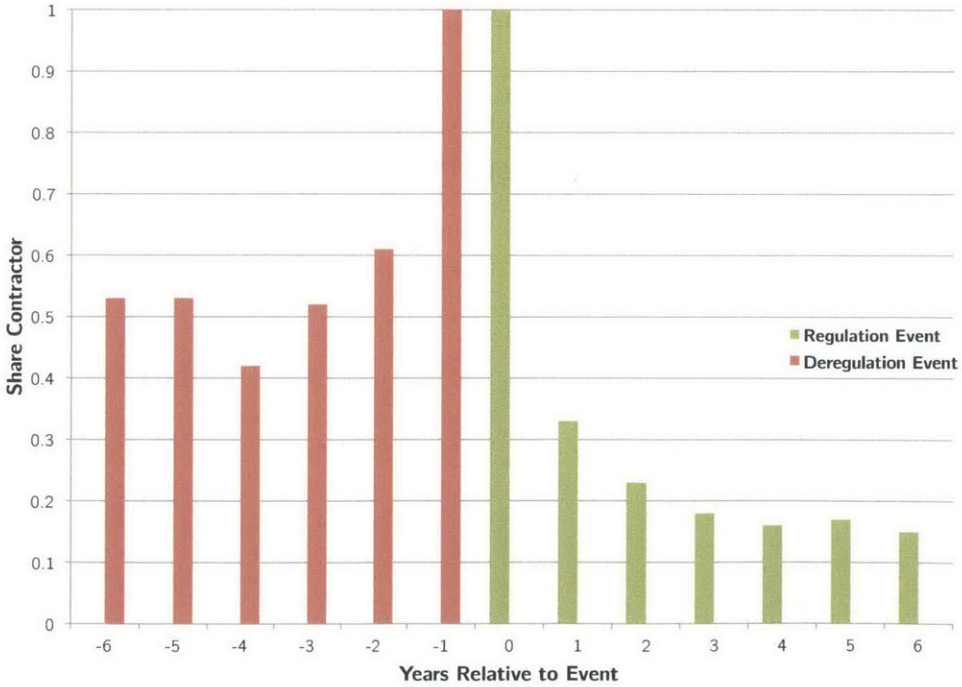

Notes: These figures graph summary statistics for sample used to construct the main event study plots presented in Figure 1.3. In Panel A I tabulate the number of establishments used to identify each lead and lag in the regulation and deregulation event studies. In Panel B I show the fraction of eventual contractors that are contractors at each node of the event studies.

Figure 1.C.2: Distribution of Contractor Spell Length

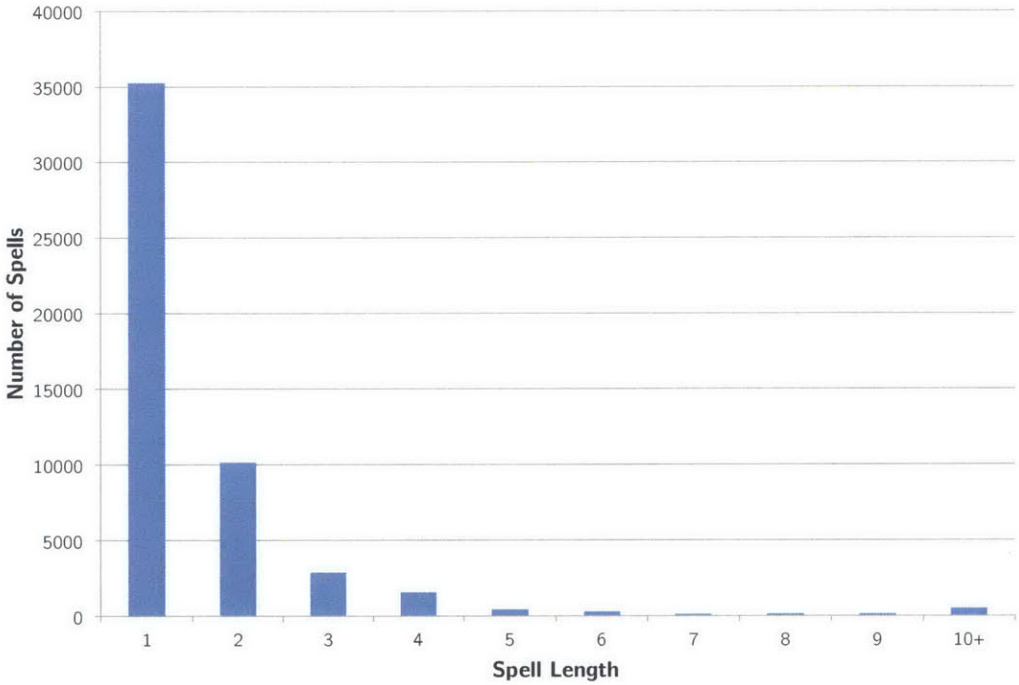

Notes: This figure plots the histogram for contractor spell length in the overlapping sample. A contractor spell is any period of consecutive years when an establishment is a contractor.

Figure 1.C.3: Number of Contractor Episodes by Establishment

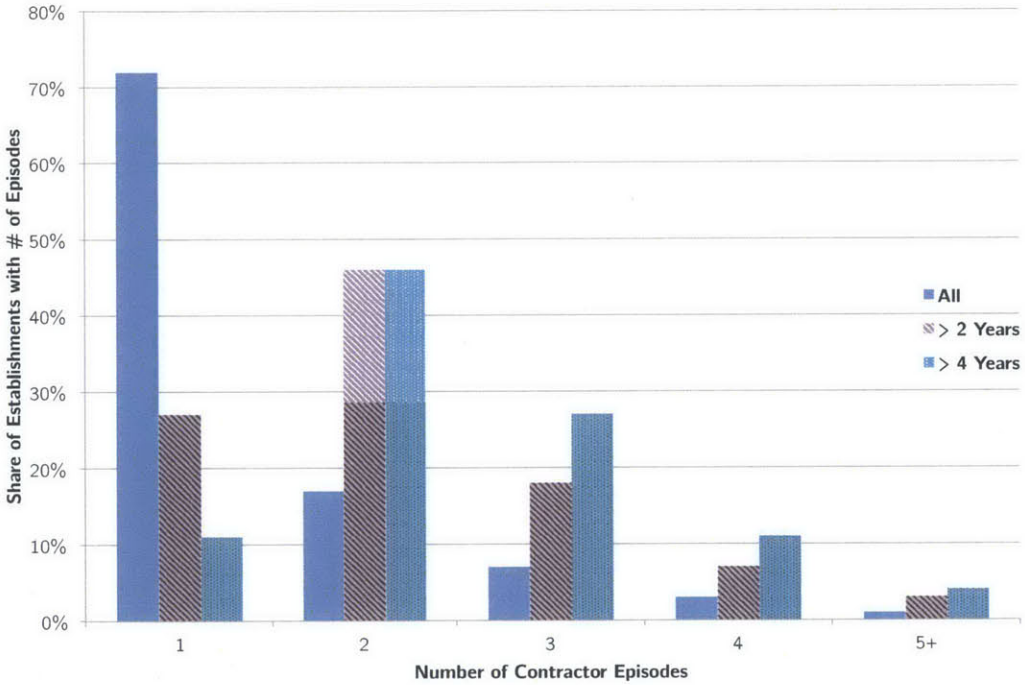

Notes: This figure plots the frequencies for the number of contractor episodes experienced by establishments in the overlapping sample. The '> 2 Years' and '> 4 Years' bars refer to eventual contractors with more than 2 or 4 years between their regulation and deregulation events.

Figure 1.C.4: Persistence and Bias

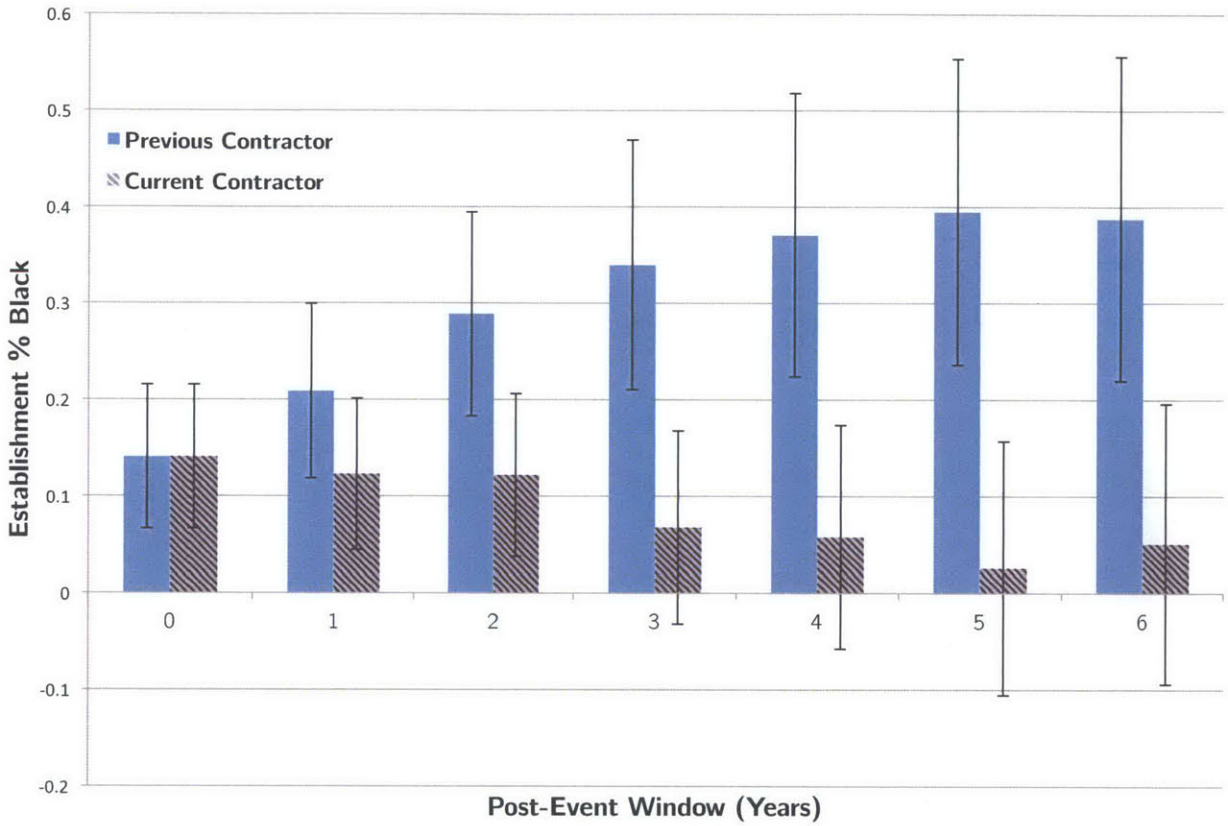

Notes: Each bar represents the coefficient estimates for a regression, along with a 95% confidence interval, with standard errors clustered at the firm level. The purple bars depict β coefficient estimates for (1.A.1). The blue bars depict $\tilde{\beta}$ coefficient estimates for (1.A.2). The estimation samples exclude establishments that enter the sample as a federal contractor. For the establishments that become contractors, I only include years of data that are at most 6 years prior to their regulation event. The 'Post-Event Window' corresponds to different estimation samples. I estimate both (1.A.1) and (1.A.2) for a series a estimation samples, moving the data window from the year of the regulation event to 6 years following the regulation event. For each sample, I also restrict the set of eventual contractors included in the estimation to those that are present for the full set of years following the regulation event. Each model includes Census division by year fixed effects, and a quadratic in log establishment size.

Table 1.C.1: EEO-1 Reporting Rates by Industry, 1990

Industry	EEO-1 Coverage Rate (%)	Industry Size (%)
Agricultural Services	22.7	0.5
Mining	90.7	0.4
Construction	11.0	5.7
Manufacturing	63.3	18.6
Transportation, Communications, Utilities	62.3	6.3
Wholesale Trade	19.4	7.2
Retail Trade	28.8	20.5
Finance, Insurance, and Real Estate	43.9	8.3
Services	30.9	32.6
Overall	37.8	100.0

Notes: Coverage rates are calculated by dividing EEO-1 reported employment by County Business Patterns employment totals for the 204 MSAs used in the analysis. Industry size is the fraction of total County Business Patterns reported employment in that industry.

Table 1.C.2: Discrimination Charge Data, Summary Statistics

	Full Data		Matched Sample	
	All Charges	Merit Resolutions	All Charges	Merit Resolutions
Total	639,058	113,782	140,124	20,247
Issue (%)				
Discharge	55	50	52	44
Discipline	9	9	10	11
Harassment	16	18	16	17
Hiring	6	5	5	5
Promotion	11	12	14	15
Suspension	4	5	5	6
Terms and Conditions	25	27	25	28
Wages	8	10	8	10
Other	24	28	24	28

Notes: Data include all Title VII discrimination charges filed with EEOC from 1990 to 2004 alleging discrimination against black workers. Matched sample includes all charges that I was able to match in the EEO-1 data, restricting to the 204 MSAs used in the analysis. Each charge can include multiple allegations, so issue percentages do not sum to 100. Merit resolutions are discrimination charges that result in a withdrawal with benefits, settlement with benefits, successful conciliation, or unsuccessful conciliation.

Chapter 2

When Work Moves: Job Suburbanization and Black Employment

2.1 Introduction

Over the last 50 years in the United States, the black unemployment rate has been roughly double the white unemployment rate (Fairlie and Sundstrom 1999). Large disparities remain after conditioning on measures of labor market skill, including education and AFQT 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% of jobs and white residents in 1950 to 30% 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 (see Holzer 1991, Ihlanfeldt and Sjoquist 1998, or Kain 1992 for older literature reviews).[1] Most recent work in this literature finds some support for spatial mismatch, though results tend to be sensitive to how job accessibility is measured.[2] 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 oppor-

[1]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).

[2]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.

tunities, where rents are typically lower.[3] Similarly, firms may choose to locate in neighborhoods with residents who are (unobservably) more productive. In this paper 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.[4] 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.[5] 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% 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% relative to white employment. Relative earnings declined by 1.1 to 2.3%. 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 black relative employment for college graduates. Notably, the estimates are driven almost entirely by job suburbanization that occurred during the 1970's. From 1970 to 1980, for every 10% decline in the fraction of MSA jobs located in the central city over this period, black relative employment rates declined by about 2.6%. From 1980 to 2000, such suburbanization is associated with only a 0.5 to 0.6% decline in black relative employment rates. Given the magnitude of the estimates over the full period, this suggests that suburbanization occurring over the 1970's 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 1940's and 50's and hence their assignment across MSAs should be orthogonal to 1970's residual labor supply shocks. I show that they appear to be unrelated to demand shocks particular to black labor

[3]Alternatively, if spatial frictions are relevant, residents who find it difficult to obtain work may sort into neighborhoods with higher employment density.

[4]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.

[5]For example, productive black households may move out of an MSA following widespread suburbanization, leaving less productive black households behind.

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 black relative employment. To the best of my knowledge, this is the first paper 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 black relative employment emerges during the 1970's 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 1970's, 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.

Though focused on spatial mismatch, this paper also relates to work on residential segregation and neighborhood effects, particularly given the potential connection between suburbanization and neighborhood composition. Neighborhood composition may be an additional mechanism through which suburbanization affects black-white labor market inequality, an issue that I explore below. Cutler and Glaeser (1997) and Ananat (2011) argue that residential segregation leads to worse labor market and social outcomes for black Americans. However, as in the existing spatial mismatch literature, these papers cannot convincingly estimate the effect of segregation on *individuals* because households can choose which city to live in. The neighborhood effects literature surmounts this problem by exploiting the (quasi-)random assignment of households to neighborhoods. Notable papers in this literature include Oreopoulos (2003), Jacob (2004), and work on the Moving to Opportunity (MTO) program (e.g. Ludwig et al. 2013). While the focus of much of this research is on the outcomes of children who grew up in these neighborhoods, Ludwig et al. (2013) find that the MTO program had no detectable effect on the economic outcomes of adult participants. By contrast, this paper studies the effects of a changing city–in particular, the changing location of employers–on the distribution of labor market outcomes among the populace. I find that job suburbanization has a substantial impact on black relative employment, but the relationship is

significantly weaker in more recent decades, including when the MTO experiment began. In future work, I will attempt to identify the forces underlying this attenuation.

In the next section, I discuss the underlying mechanisms that may produce spatial mismatch in more detail. In section 2.3 I discuss the data, how I construct the synthetic panel, and how I measure job suburbanization. In section 2.4 I discuss the empirical strategy, present results, and discuss interpretation. In section 2.5 I describe the Equal Employment Opportunity Commission data and explore whether my results may be driven by job displacements. Section 2.6 concludes.

2.2 Mechanisms of Spatial Mismatch

The spatial mismatch hypothesis predicts that the suburbanization of work, combined with the concentration of black households in the central city, reduces black relative employment. Coulson et al. (2001) develop a general equilibrium spatial search model that identifies the conditions necessary to generate higher unemployment rates in the central city due to spatial frictions. The model can be easily adapted to identify the conditions under which job suburbanization reduces black employment.

In the Coulson et al. framework, the following conditions are sufficient for spatial mismatch to emerge in equilibrium (Johnson 2006):

1. Black households are relatively constrained to residing in the central city,

2. Commuting or job search costs (as a function of distance) are nontrivial,

3. Firms face higher (non-labor) costs (either fixed or production costs) in the central city than in the suburbs.

Conditions (1) and (2) are essential because black labor must find it relatively difficult to work in the suburbs. With no search or commuting costs, the spatial distribution of work would have no bearing on labor market outcomes. Even with search and commuting costs, if black and white households find it equally costly to reside in the suburbs, black households would be equally able to follow work to the suburbs. Indeed, there is evidence that black households faced relatively high barriers to suburban residence over this period due to discrimination in housing and mortgage markets (Yinger 1986, 1995). In addition, given their low levels of wealth (Blau and Graham 1990; Barksy et al. 2002), black households may be more likely to face binding liquidity constraints in securing suburban housing. Also note that, at least in the short run, black households may be relatively constrained to the central city simply because that is where they are initially concentrated and moving costs are nontrivial.

Even with constraints on residential mobility and nontrivial search and commuting costs, the mobility of firms will tend to equalize employment opportunities across the central city and suburbs. If productive workers remain in the central city, competitive forces should induce the entry of a sufficient number of firms to absorb that labor until subsequent entry remains unprofitable. For

example, suppose job suburbanization was entirely driven by the residential suburbanization of employees (e.g. due to relative improvements in suburban amenities), a commonly cited cause (Glaeser and Kahn 2001). In this case, it is not clear why the movement of firms to the suburbs would have any detrimental effect on labor that is constrained to the central city. There would be fewer jobs in the central city; however, if the residual labor remains equally productive, they may remain equally attractive to potential employers. As emphasized in Coulson et al. (2001), for the predictions of the spatial mismatch hypothesis to hold, it is critical that job suburbanization be in part driven by declines in the relative costs of operating in the suburbs faced by firms. This makes suburban entry more attractive, even ignoring concerns about labor accessibility. Holding central city labor productivity constant, central city entry becomes relatively less profitable, increasing the relative unemployment rate in the central city.

Indeed, there is evidence that a number of non-labor factors induced firm entry in the suburbs by reducing the relative fixed and production costs associated with operating in the suburbs. These include innovations in transportation and transportation infrastructure, lower suburban land costs, and the formation of agglomeration economics in the suburbs (Glaeser and Kahn 2001).

In addition, labor suburbanization can also reduce the effective productivity of workers that remain in the central city if there are complementarities between types of labor in the production function, production exhibits increasing returns to scale, or product demand is local. For example, if the presence of skilled managers in a plant increases the productivity of unskilled workers, then the migration of skilled labor out of the central city may reduce the effective productivity of the unskilled workers left behind.

In the Coulson et al. framework, there are three factors that can induce job suburbanization: (1) declines in the relative non-labor costs of operating in the suburbs; (2) the suburbanization of labor; and (3) shocks to productivity that make labor that remains in the central city becomes relatively less productive. Only suburbanization spurred by (1) and (3) can produce an equilibrium negative correlation between job suburbanization and black relative employment. Under (3), black labor can become relatively less productive, and because black households are concentrated in the central city, this would induce firms to shift to the suburbs. Only job suburbanization produced by (1) can *cause* declines in black relative employment. This relationship is a direct implication of the spatial mismatch hypothesis. I will test the spatial mismatch hypothesis by estimating the relationship between job suburbanization and black relative employment while accounting for the role unobserved productivity may play in this relationship.

Though outside the Coulson et al. model, suburbanization can also alter racial inequality by changing the composition of residential neighborhoods. If residential and job suburbanization occur in tandem and the rate of household decentralization varies with race or income, then suburbanization may increase residential segregation by race or income. Residential segregation may affect labor market inequality directly through channels that are orthogonal to spatial mismatch (Cutler and Glaeser 1998). In the analysis below I assess the potential role for neighborhood compositional changes in explaining the results.

2.3 Data and Implementation

The ideal data source for this study would include longitudinal information on a large sample of individuals' labor market outcomes and location of residence, as well as the spatial distribution of work establishments. Lacking such a data set, I use Census data from 1970, 1980, 1990, and 2000. I focus on these Census years for two reasons. First, the second wave of the Great Migration ends around 1970. Analysis of Census data from earlier than 1970 would be complicated by the large and potentially endogenous black migration flows over this period. Second, MSA is not identified in the 1960 IPUMS Census data.

In analyzing labor market outcomes, I restrict attention to men and women between the ages of 24 and 64 who are non-Hispanic white or black. To measure employment I use an indicator for whether an individual is currently 'active', meaning employed or in school.[6] In practice, this measure reflects employment because only a small fraction of individuals in my sample report being in school and this does not differ significantly by race.[7] For this reason, I refer to 'active' and 'employed' interchangeably. I restrict the analysis to the 71 metropolitan statistical areas (MSAs) with the largest black populations as defined in 1970.

Combining Integrated Public Use Microdata (IPUMS) from each Census, I construct a synthetic panel by collapsing individuals into cells and merging cells across years. The 1970 Census data is a 2% national samples, while the remaining years are 5% national samples. I group individuals by MSA of residence, gender, race, education, and cohort. I exclude the institutionalized because individuals in that population may not be residing in their relevant labor market. This may attenuate the relevant estimates below given that incarceration rates began to increase substantially over this period and a large share of black adults was incarcerated (Western and Pettit 2000). Patterns for women should be less susceptible to this issue. I divide the sample into 3 education groups: less than high school graduate, high school graduate, and college graduate. I also divide the sample into 3 cohorts, those who are the following ages in 1970: 24-33, 34-43, and 44-54. I group by cohort rather than age because the intention is to follow the same group of individuals from 1970 to 1980. This is an important aspect of the analysis relative to the existing, primarily cross-sectional literature because households may sort across MSAs in a manner related to unobserved determinants of employment status. Of course, the composition of cells may change from decade to decade due to migration. I explore the role migration plays in the analysis in section 2.4.3.

Summary statistics on the demographics and labor market outcomes for the synthetic cohorts are presented in Table 2.1. Black men are less active in the labor market than white men in 1970 and experience larger proportional declines in employment from 1970 to 1980. By contrast, black women are more active in the labor market than white women in 1970. However, the white women experience larger across- and within-cell increases in employment rates from 1970 to 1980.

To measure job suburbanization, I use various Census data products. Measuring the spatial

[6]Results are similar if I use weeks worked as the employment measures.

[7]This measure follows Cutler and Glaeser (1997) who, in studying the effects of residential segregation on youth, define 'idle' as not employed and not in school.

Table 2.1: Sample Descriptive Statistics, Cell-Level

	Black				White			
	1970	1980	1990	2000	1970	1980	1990	2000
Share	0.131	0.150	0.145	0.165	0.869	0.850	0.855	0.835
Female	0.562	0.562	0.568	0.569	0.512	0.513	0.511	0.510
1936-1945	0.279	0.168	—	—	0.340	0.229	—	—
1946-1955	0.339	0.200	0.161	—	0.314	0.207	0.189	—
1956-1965	0.382	0.249	0.192	0.225	0.346	0.232	0.205	0.253
1966-1975	—	0.382	0.301	0.346	—	0.333	0.298	0.359
1976-1985	—	—	0.346	0.429	—	—	0.308	0.387
<HS Grad	0.558	0.381	0.272	0.211	0.291	0.194	0.119	0.085
HS Grad	0.398	0.528	0.593	0.621	0.540	0.572	0.578	0.566
College Grad	0.045	0.091	0.135	0.168	0.168	0.234	0.304	0.349
Active, Male	0.869	0.767	0.758	0.701	0.947	0.888	0.887	0.853
	(0.061)	(0.121)	(0.141)	(0.143)	(0.030)	(0.096)	(0.105)	(0.107)
Active, Female	0.571	0.627	0.693	0.664	0.480	0.595	0.713	0.707
	(0.120)	(0.161)	(0.170)	(0.160)	(0.081)	(0.132)	(0.141)	(0.119)
π^{cc}	0.535	0.503	0.472	0.435	0.526	0.497	0.461	0.424
	(0.120)	(0.138)	(0.167)	(0.181)	(0.117)	(0.127)	(0.151)	(0.165)
$\Delta \log(\pi^{cc})$	—	-0.084	-0.100	-0.106	—	-0.080	-0.097	-0.105
	—	(0.114)	(0.137)	(0.104)	—	(0.108)	(0.128)	(0.086)
$\pi^{cc}_{res,race}$	0.795	0.734	0.678	0.617	0.393	0.324	0.292	0.270
	(0.131)	(0.155)	(0.181)	(0.207)	(0.157)	(0.150)	(0.153)	(0.155)

Notes: Includes 71 MSAs with largest black populations in 1970. Weighted by cell size. π^{cc} refers to the fraction of MSA jobs located in the central city. $\pi^{cc}_{res,race}$ refers to the fraction of black or white MSA residents living in the central city.

distribution of work consistently across years is complicated by the fact that central city definitions change significantly with the 1990 Census. In particular, many cities that are defined as suburbs in 1980 are classified as central cities in 1990. These changes make it difficult to construct consistent measures of job suburbanization after 1980 using only Census IPUMS data. Instead, I combine IPUMS data with tabulations from the Census Transportation Planning Package (CTPP) to measure the spatial distribution of work in 1990 and 2000. The CTTP data include tabulations reporting the total number of individuals working at various levels of geography. I divide those totals into central city and suburb using 1980 Census definitions for central cities. For 1970 and 1980, I use the Census IPUMS data used above. In the Census micro data, I measure the spatial distribution of work using the Census indicator for whether an individual works in the central city or outside the central city (suburbs) of an MSA.

2.3.1 Measuring Job Suburbanization

The ideal source of across-MSA variation in job suburbanization would come solely from the locational choices of entering (or moving) firms. In particular, suppose every MSA began with the same types of firms (e.g. as defined by industry and size) and initial central city-suburban division of work. Moreover, suppose the same types of firms exited and entered each MSA market at the same rates from 1970 to 2000, and that the spatial distribution of exiting firms by type was also similar across MSAs. The only variation across MSAs in 2000 would derive from the location choices of entering (or moving) firms and hence changes in employment rates across groups could be attributed to changes in spatial distribution of work rather than other confounding factors (e.g. changes in sector composition, overall growth, or turnover).

In practice, variation in job suburbanization deviates from the ideal in primarily three ways. First, MSAs differ in their initial spatial distribution of work. My measure of job suburbanization must account for these differences in initial conditions. Second, there is across-MSA variation in the types of firms that are present in 1970 and those that exit and enter by 2000. This may be important because changes in the composition of jobs that occur concurrently with job suburbanization may also independently determine group employment rates. Third, the rate of exit and entry and the location of exiting firms vary across MSAs. This presents two issues: (a) overall MSA job growth may be correlated with job suburbanization and determine group employment rates; (b) if job suburbanization is in part driven by differences in exit rates, the relationships between job suburbanization and group employment rates may be driven by heterogeneity in exposure to job displacements across groups within MSAs or across-group heterogeneity in the effects of job displacement on subsequent employment rates. I discuss how I measure job suburbanization and account for these deviations in the rest of this section except for point 3(b), which I discuss in section 2.5.

To analyze the effects of job suburbanization using variation across MSAs, I first need a measure of job suburbanization that can be applied consistently across MSAs with differing initial spatial distributions of work. I focus on the percentage change in the number of central city jobs due

Table 2.2: Job Suburbanization By City

Rank	MSA Name	$\Delta \log(\pi^{cc})$
1970-1980		
1.	Detroit, MI	-0.323
2.	Atlanta, GA	-0.305
3.	Nashville, TN	-0.295
4.	Minneapolis-St. Paul, MN	-0.265
5.	Tampa-St. Petersburg-Clearwater, FL	-0.186
...	...	...
33.	Los Angeles-Long Beach, CA	-0.056
34.	Akron, OH	-0.052
35.	Hartford-Bristol-Middleton-New Britain, CT	-0.048
36.	Knoxville, TN	-0.047
37.	Denver-Boulder, CO	-0.046
...	...	...
67.	West Palm Beach, FL	0.206
68.	New Haven-Meriden, CT	0.219
69.	Omaha, NE/IA	0.223
70.	Miami-Hialeah, FL	0.229
71.	Charleston-N. Charleston, SC	0.402
1970-2000		
1.	Detroit, MI	-0.804
2.	Atlanta, GA	-0.742
3.	Richmond-Petersburg, VA	-0.709
4.	St. Louis, MO-IL	-0.645
5.	Wilmington, DE/NJ/MD	-0.575
...	...	...
33.	Phoenix, AZ	-0.285
34.	San Francisco-Oakland-Vallejo, CA	-0.270
35.	Kansas City, MO-KS	-0.217
36.	Hartford-Bristol-Middleton-New Britain, CT	-0.216
37.	Riverside-San Bernardino, CA	-0.199
...	...	...
67.	Omaha, NE/IA	0.048
68.	Bridgeport, CT	0.093
69.	Norfolk-VA Beach-Newport News, VA	0.100
70.	Chattanooga, TN/GA	0.101
71.	West Palm Beach, FL	0.106

Notes: Includes 71 MSAs with largest black populations in 1970.

to the change in the spatial distribution of work, what I term the *location effect*. The number of central city jobs may change because the whole MSA is growing or shrinking—the *scale* effect—or via the location effect. More concretely, let T_t denote the number of jobs in an MSA at time t ($t = 70, 80, 90, 00$); let π_t^{cc} denote the fraction of MSA jobs located in the central city; and let T_t^{cc} denote the number of central city jobs. The change in the log of the number of central city jobs can be decomposed as follows:

$$\begin{aligned}\log T_{t_1}^{cc} - \log T_{t_0}^{cc} &= \log(\pi_{t_1}^{cc} T_{t_1}) - \log(\pi_{t_0}^{cc} T_{t_0}) \\ &= \underbrace{[\log \pi_{t_1}^{cc} - \log \pi_{t_0}^{cc}]}_{\textit{location}\text{ effect}} + \underbrace{[\log T_{t_1} - \log T_{t_0}]}_{\textit{scale}\text{ effect}}.\end{aligned}$$

Hence, I measure job suburbanization using $\Delta_{t_1,t_0} \log \pi^{cc} = \log \pi_{t_1}^{cc} - \log \pi_{t_0}^{cc}$. Table 2.1 also presents summary statistics on the job suburbanization faced by the synthetic cohorts. In each decade, π^{cc} decreases by 8-10%.

The above is an MSA-level measure of job suburbanization; however, the analysis is at the cell level. For some analyses, I construct cell-level measures of exposure to job suburbanization by using additional information on the industries and occupations that individuals in different cells tend to work in. Doing this serves two purposes. First, cell-specific measures may provide more accurate measures of a cell's exposure to job suburbanization. For example, if men are concentrated in manufacturing relative to women, the suburbanization of the manufacturing sector is more relevant in measuring the exposure of men to job suburbanization. Second, by constructing cell-specific measures of job growth analogously, I can examine whether the relationship between job suburbanization and black employment is driven by growth or declines in industries that blacks tend to work in.

In particular, I first divide the approximately 300 detailed industries defined in the Census into 15 major industry groups.[8] I then compute $\Delta \log \pi^{cc}$ separately for each major sector within each MSA. To form a measure of job suburbanization for each cell, I take a weighted average of the decentralization of sectors in each MSA, where the weights are the 1970 shares of the cell's employed members working in each sector.[9] More concretely, let α_{s,t_0}^{mg} denote the share of group g (race by education by gender by cohort combination), MSA m members employed in year t_0 who work in sector s, so that $\sum_s \alpha_{s,t_0}^{mg} = 1$. I define $\bar{\Delta}_{mg} \log(\pi^{cc})$ as

$$\bar{\Delta}_{mg} \log(\pi^{cc}) = \sum_s \alpha_{s,t_0}^{mg} \Delta \log(\pi_{s,t_0}^{cc}).$$

I also construct analogous measures based on occupations, dividing about 300 detailed occupations

[8]These major groups consist of: agriculture, forestry, and fisheries; mining; construction; nondurable goods manufacturing; durable goods manufacturing; transportation, communications, and other public utilities; wholesale trade; retail trade; finance, insurance, and real estate; business and repair services; personal, entertainment, and recreation services; professional services, health; professional services, education; professional services, other; and public administration.

[9]This strategy is similar to the method used in Card (2001) to measure the substitutability of immigrants and natives using their distributions across occupations.

into 10 major categories.

This measure of job suburbanization may introduce error in (at least) three ways. First, the distinction between central city and suburb is somewhat artificial in describing the changes in the spatial distribution of work. Second, the weights I use across sectors (industry or occupation) may be inappropriate or the sectors I use may be an inappropriate partition of jobs. Reassuringly, my results are unchanged if I repeat the analysis using a finer partition of industries or occupations. Third, the sector shares α will change from year to year, either due to changes in marginal distribution of cells across sectors or changes in the overall composition of sectors.

An additional potential complication derives from the fact that the suburbanization and growth (or decline) of MSAs or sectors may happen simultaneously, or one may cause the other. For example, suburbanization may occur as an artifact of an MSA's or industry's decline if those declines tend to begin in the central city, or suburbanization may lead to the decline of certain industries. However, the exercise calls for estimating the relationship between suburbanization and changes in employment rates *conditional* on job growth, changes in sectoral composition, and other changes in the local economy that may determine employment rates.

Fortunately, suburbanization and job growth are uncorrelated at the MSA and MSA by sector level. At the MSA by industry level, the (weighted) correlation between suburbanization and growth is a negligible 0.029. These results imply that I am unlikely to confound the effects suburbanization with growth or changes in sector composition.

Still, in some specifications, I include measures MSA or cell-specific demand as in Bartik (1991) as controls. This demand index is a prediction of employment growth derived by interacting cross-sectional differences in sector composition with notional changes in sector employment shares. Formally, the index is defined as

$$\theta^{mg}_{i,t_1,t_0} = \sum_{s}^{S} \alpha^{mg}_{s,t_0} \frac{\alpha^{-m}_{s,t_1} - \alpha^{-m}_{s,t_0}}{\alpha^{-m}_{s,t_0}}$$

where $\alpha^{-m}_{s,t}$ is the share of national employment outside of MSA m in sector s at year t. Unsurprisingly, the inclusion of these demand shock controls has little impact on the results.

Descriptive statistics on the spatial distributions of and share of workers in different sectors are presented in the Appendix. From 1970 to 1980 the fraction of work in the central city declines by an average of 7.4% across MSAs. However, this average masks substantial heterogeneity across MSAs and within MSAs across sectors. The weighted mean value of $\bar{\Delta}_{mg} \log(\pi^{cc})$ across cells is -0.095 for blacks with standard deviation 0.118 and -0.094 for whites with standard deviation 0.113.

2.4 Empirical Strategy and Results

I test the spatial mismatch hypothesis by estimating the relationship between job suburbanization and black relative employment. As discussed in section 2.2, in the presence of spatial mismatch, job suburbanization will have a weakly negative correlation with black relative employment, conditional on worker productivity. The magnitude of this relationship will depend on the significance of spatial

frictions as well as the factors contributing to job suburbanization. In particular, in the Coulson et al. framework, spatial mismatch will only generate black employment declines to the extent that job suburbanization is driven by declines in the relative costs of suburban entry and operating costs faced by firms. Unfortunately, I cannot observe these entry and operating costs directly. Instead, I assume that some portion of observed job suburbanization is due to changes in relative costs and estimate the relationship between job suburbanization and black relative employment while attempting to control for any changes in worker productivity.

To determine the empirical relationship between black relative employment and job suburbanization, I estimate linear models of the following form, where all coefficients are allowed to differ by race:

$$\Delta(emp)_{mg} = \sum \alpha^g + \beta \Delta \log(\pi_m^{cc}) + f(emp_{mg}^{t_0})\gamma + X_m\delta + \epsilon_{mg} \tag{2.1}$$

where α^g are group level fixed effects. Some specifications include a vector X, a set of additional MSA-level controls including regional indicators and the black share of the MSA population at baseline. As my measure of employment, I use the fraction of the cell that is active in the labor market, that is, either employed or in school. I estimate models primarily in logs.[10] In addition, I look at log average annual earnings. In some specifications I do not condition on baseline employment, but in others I specify $f(\cdot)$ as a quadratic function. I include a control for a polynomial in baseline employment because employment growth may depend nonlinearly on baseline employment. For example, because the cell employment rate is capped at 1, cells with high baseline employment have very little potential for growth relative to cells with lower baseline employment. In addition, conditioning on baseline employment should mitigate the potential for bias introduced by mean reversion. The coefficient β_1 has the following interpretation: a 10% decline in the fraction of MSA jobs located in the central city is associated with at $.10 \times \beta_1\%$ (or percentage point when the outcome is in levels) decline in cell employment.

Before describing the baseline results, I explore how job suburbanization relates to other baseline cell-level and MSA-level characteristics. Table 2.1 provides coefficient estimates for models in the form of (2.1) except the left hand side variable is replaced by various group-level and MSA-level covariates other than employment growth. I relate these correlates to short run (1970 to 1980) and long run (1970 to 2000) subsequent job suburbanization. One concern with my empirical strategy is that job suburbanization may occur in areas where employment is already declining, particularly for black labor. Though the 1960 Census does not identify MSA, for half the respondents in the 1970 Census, I can observe their employment status in 1965. To measure pre-trends by cell, I use the change in employment rates by cell, assigning individuals to MSAs using their residence in 1970. I find that cells that experienced more job suburbanization from 1970 to 1980 had somewhat lower employment growth from 1965 to 1970, but the magnitude is small. Moreover, the difference in trend between black and white cells is statistically insignificant. I also find that job suburbaniza-

[10] Results are very similar if I use log weeks worked.

Table 2.1: Job Suburbanization and Group and MSA Characteristics, 1970

Dependent Variable	1970-1980 $\Delta \log(\pi^{cc})$	1970-1980 $\Delta \log(\pi^{cc})$ × black	1970-2000 $\Delta \log(\pi^{cc})$	1970-2000 $\Delta \log(\pi^{cc})$ × black	N MSAs
Changes over 1965-1970					
Δ Fraction Active, Group	0.038*	0.024	0.020	-0.009	71
	(0.019)	0.039	(0.027)	(0.015)	
Δ Log Fraction Active, Group	0.067~	0.044	0.036	-0.030	71
	(0.038)	0.033	(0.052)	(0.024)	
Baseline (1970) Characteristics					
Fraction Active, Group	0.022	0.044	0.010	-0.060*	71
	(0.028)	0.040	(0.021)	(0.022)	
Log Earnings, Group	-0.165*	-0.386**	0.088	-0.382**	71
	(0.070)	(0.132)	(0.153)	(0.129)	
Fraction of Work in Central City, Group	0.092	0.143~	0.166*	0.022	71
	(0.157)	(0.079)	(0.063)	(0.036)	
Fraction Black, MSA	-0.066	0.059	-0.007	-0.010	71
	(0.069)	(0.060)	(0.038)	(0.036)	
Fraction Black, Central City	-0.366*	-0.058	-0.248**	-0.045	69
	(0.109)	(0.112)	(0.080)	(0.064)	
Violent Crime Rate, MSA	-1.716	0.092	0.069	-0.243	70
	(1.418)	(0.653)	(1.011)	(0.289)	
Property Crime Rate, MSA	-0.516	-0.191	-0.593	-0.112	70
	(1.197)	(0.374)	(0.409)	(0.115)	
Dissimilarity Index, MSA	-0.193*	-0.042	-0.129**	-0.022	68
	(0.079)	(0.038)	(0.036)	(0.012)	

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

tion does not predict the *level* of baseline cell employment rates for blacks or whites. However, job suburbanization is positively correlated with baseline earnings, particularly for blacks. As a robustness check, in some specifications I include baseline earnings as a control. I also find that job suburbanization is uncorrelated with the baseline black share of the MSA population. In contrast, job suburbanization is a strong predictor for the baseline black share of the *central city*. This is consistent with research suggesting that black in-migration to central cities was a major cause of 'white flight' (Boustan 2010).

I also relate job suburbanization to baseline MSA crime rates. Data on reported crimes come from the FBI's Uniform Crime Reports. The UCR reports crimes per 100,000 for 7 types of offenses: murder, rape, robbery, aggravated assault, burglary, larceny, and motor vehicle theft. I divide these 7 offenses into two categories violent and property crimes—sum within these categories and standardize these sums to have mean zero and standard deviation one across MSAs. Job suburbanization does not have a statistically significant relationship with either baseline violent or property crime rates; however, these estimates are noisy.

Table 2.2: Job Suburbanization and Labor Market Outcomes, By Decade, 1970-2000

	1970-1980		1980-1990		1990-2000	
Outcome: Δ log (Employment Rate)	(1)	(2)	(3)	(4)	(5)	(6)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	-0.036	-0.025	0.017	0.009	-0.023	-0.030
	(0.046)	(0.052)	(0.040)	(0.028)	(0.021)	(0.018)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.242**	0.277**	0.050	0.060	-0.028	0.013
	(0.059)	(0.069)	(0.039)	(0.047)	(0.040)	(0.026)
Outcome: Δ log (Earnings)	(7)	(8)	(9)	(10)	(11)	(12)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	-0.030	-0.053	0.204	0.195	-0.043	-0.020
	(0.059)	(0.068)	(0.149)	(0.158)	(0.037)	(0.045)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.370**	0.157~	0.084	0.032	-0.061	-0.030
	(0.090)	0.086	(0.076)	(0.082)	(0.088)	(0.053)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic in Baseline	No	Yes	Yes	No	Yes	Yes
N Cells	1944	1944	2315	2315	2399	2399
N MSAs	71	71	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

Finally, I look at residential segregation and suburbanization to gauge the potential role for changing neighborhood composition in producing my results. I find that job suburbanization has a significant positive relationship with baseline residential segregation and segregation growth, where the significance is marginal for the latter outcome. However, the magnitude is modest: a 10% decline in the fraction of MSA jobs located in the central city is associated with a 0.01-0.015 point increase in the dissimilarity index, which estimates from Collins and Margo (2000) imply would predict a 0.2 to 0.3 percentage point relative increase in the idle rate for black 20-30 year olds.[11]

2.4.1 Baseline Estimates

In estimating (2.1), I restrict the analysis to cells with at least 25 observations, and weight cells by their size. This leaves 1944 cells over 71 MSAs. I cluster standard errors at the MSA level. I first estimate separate models for each decade. Table 2.2 provides these baseline estimates. In the top panel, the outcome is log employment rate. In the bottom panel, the outcome is log average income. Even columns include controls for second order polynomial in baseline employment.

Across specifications and decades, changes in the spatial distribution of work have little association with white employment growth. The coefficient is generally small in magnitude and statistically indistinguishable from zero at the 10% level. Over the 1970's, job suburbanization is associated with declines in black employment, and the relationship is highly statistically significant. The coefficients are similar whether or not I condition on baseline employment. The β^B coefficient of 0.242 to 0.277 implies that a 10% decline in the fraction of MSA jobs located in the central city—about the mean

[11] I use the coefficient estimates from Collins and Margo (2000) for 1980 using all available cities.

Table 2.3: Job Suburbanization and Labor Market Outcomes in Longer Run

	1970-1990		1980-2000		1970-2000	
Outcome: Δ log (Employment Rate)	(1)	(2)	(3)	(4)	(5)	(6)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	-0.029	-0.004	-0.023	-0.028	-0.043	-0.023
	(0.070)	(0.047)	(0.036)	(0.017)	(0.045)	(0.033)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.238**	0.181**	0.050*	0.057**	0.210**	0.141**
	(0.049)	(0.055)	(0.022)	(0.021)	(0.042)	(0.038)
Outcome: Δ log (Earnings)	(7)	(8)	(9)	(10)	(11)	(12)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	0.074	0.086	0.126	0.118	0.021	0.030
	(0.098)	(0.106)	(0.098)	(0.104)	(0.060)	(0.067)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.239**	0.144~	0.031	-0.019	0.227**	0.114*
	(0.075)	(0.074)	(0.030)	(0.033)	(0.064)	(0.048)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic in Baseline	No	Yes	No	Yes	No	Yes
N Cells	1327	1327	1588	1588	686	686
N MSAs	71	71	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

level experienced over this period—is associated with about a 2.4% to 2.8% decline in black cell fraction employed, relative to white employment. The estimates for annual earnings have the same sign, but the coefficient is more dependent on the inclusion of baseline earnings as a control.

By contrast, the β^B coefficients for log employment in later decades are significantly smaller in magnitude and not statistically distinguishable from zero. For the 1980's, the coefficient ranges from 0.050 to 0.060, and for the 1990's the coefficient estimates are roughly centered at zero. The relationship between job suburbanization and black relative employment appears to fade out with each passing decade. The pattern of coefficients is similar for log earnings.

Next, I estimate analogous models over two decades (1970 to 1990 and 1980 to 2000), and over the full period (1970 to 2000). Long run responses may differ if, for example, labor market adjustments occur slowly. The results are presented in Table 2.3. Again, the outcome is log employment rate in the top panel and log average income in the bottom panel. From 1970 to 1990, a 10% decline in the fraction of jobs located in the central city is associated with a 1.8% to 2.4% decrease in black relative employment. From 1980 to 2000, this reduces to 0.5% to 0.6%, though the estimates remain statistically significant at the 5% level. Over the full period, a 10% decline in the fraction of jobs located in the central city is associated with a 1.4% to 2.1% decrease in black relative employment. This relationship is estimated using a single cohort of workers: those who were 25-34 in 1970, and 55-64 in 2000. Figure 2.1 displays this relationship graphically, plotting $\Delta_m \log(\pi^{cc})$ against changes in log black and white employment rates for these cells. The combination of results in Tables 2.2 and 2.3 suggest that while the associated between job suburbanization and black employment has weakened over time, the relationship is largely persistent.

To assess the robustness of these estimates, I re-estimate (2.1) including additional MSA-level

Figure 2.1: Job Suburbanization and Changes in Employment Rates, 1970-2000

(A) Synthetic Panel, Black

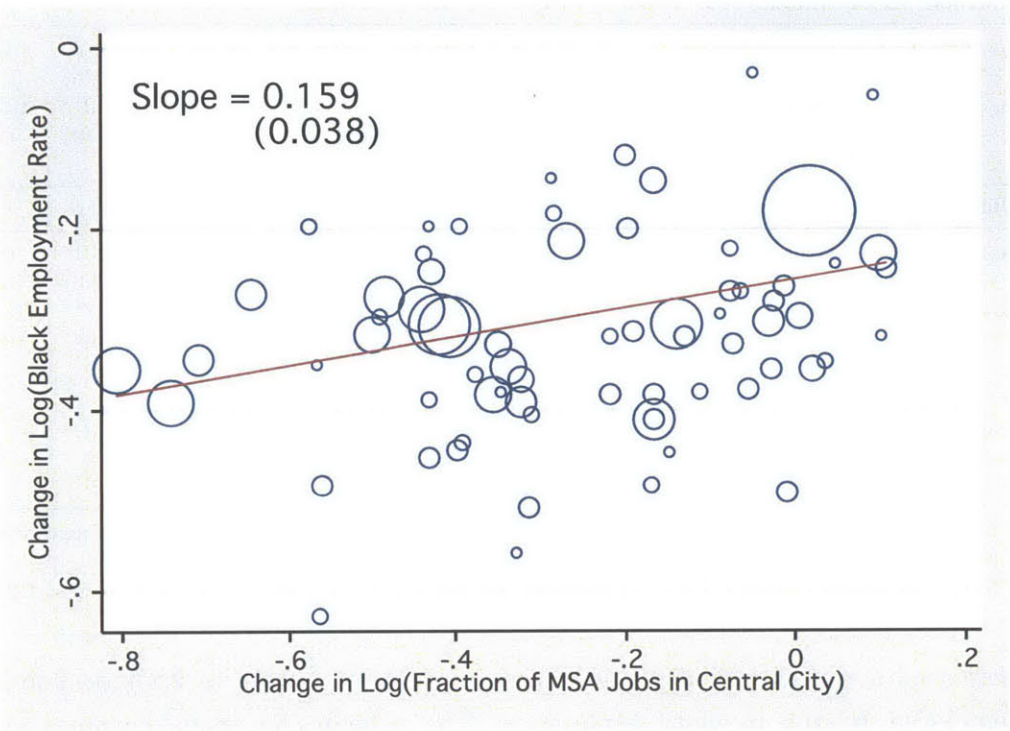

(B) Synthetic Panel, White

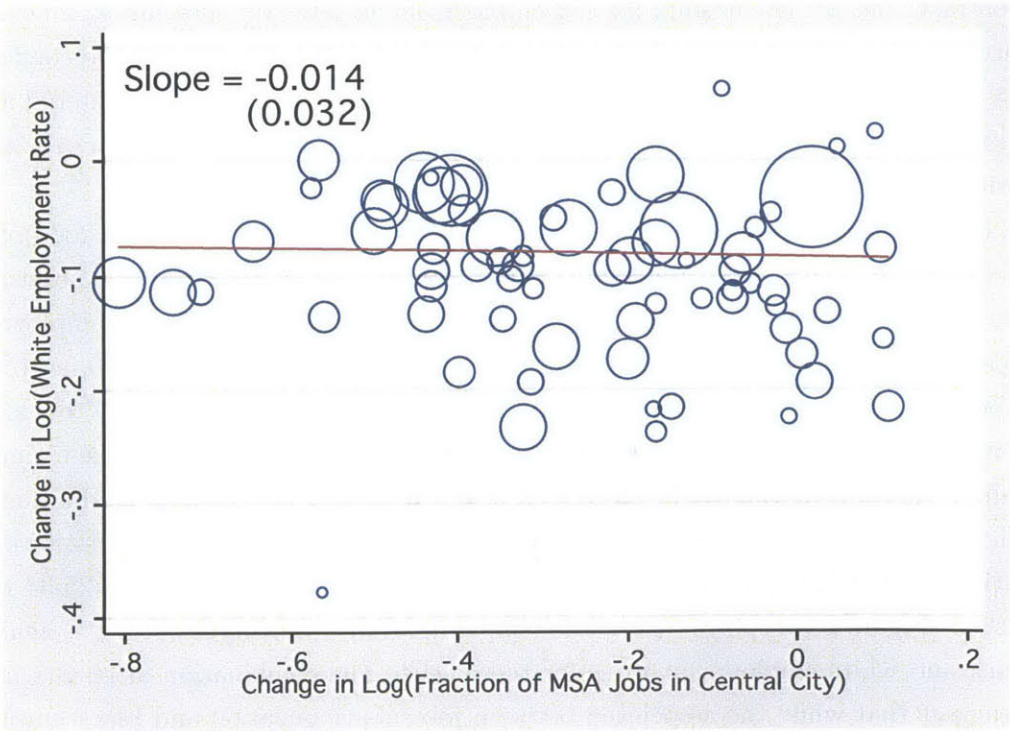

Notes: This figure plots changes in black and white employment rates against job centralization across MSAs from 1970 to 2000. See Section 2.4.1 for details on the construction of synthetic cohorts.

Table 2.4: Job Suburbanization and Labor Market Outcomes, By Group, 1970-1980

Outcome: Δ log (Emp, Rate)	HS Dropout (1)	HS Grad (2)	College Grad (3)	Men (4)	Women (5)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	0.026	-0.037	-0.054	-0.020	-0.021
	(0.070)	(0.052)	(0.038)	(0.033)	(0.070)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.265**	0.232**	0.003	0.277**	0.265**
	(0.078)	(0.058)	(0.045)	(0.078)	(0.084)
Outcome: Δ log (Earnings)	(6)	(7)	(8)	(9)	(10)
$\bar{\Delta}_{mg} \log(\pi^{cc})$	-0.041	-0.021	-0.099	-0.049	-0.054
	(0.082)	(0.070)	(0.063)	(0.080)	(0.071)
$\bar{\Delta}_{mg} \log(\pi^{cc}) \times$ black	0.182*	0.089	0.020	0.030	0.222*
	(0.087)	(0.106)	(0.249)	(0.099)	(0.100)
Group FEs	Yes	Yes	Yes	Yes	Yes
Quadratic in Baseline	Yes	Yes	Yes	Yes	Yes
N Cells	780	695	469	2315	2399
N MSAs	71	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

controls for regional fixed effects and the initial black share of the MSA population. I focus on the 1970-1980 period, when the relationship between suburbanization and black employment is strongest. The results are presented in the Appendix in Table 2.A.1. The general results are unchanged; the coefficient estimates for $\Delta \log(\pi^{cc}) \times$ *black* are slighted muted and estimated less precisely, though still statistically significant at the 1% or 5% level. This remains true when I also include a quadratic in baseline log earnings as a control.

Finally, I examine heterogeneity in the baseline estimates by education and gender. These estimates are presented in Table 2.4. Again I focus on the 1970-1980 period. All models include a quadratic in baseline employment or earnings. For employment, the relationship is equally present for high school dropouts and high school graduates, and men and women. For these groups, β^B ranges from 0.232 to 0.277. By contrast, there is no correlation between job suburbanization and black relative employment for college graduates. For earnings, the relationship is present for women but not men, and weaker for high school graduates than it is high school dropouts. Again, there is no relationship for college graduates.

2.4.2 Sector-Specific Suburbanization and Demand Shocks

For reasons discussed in section 2.3.1, one concern with interpreting the baseline estimates is that my measure of job suburbanization is MSA level. In this section, I replace $\Delta \log(\pi^{cc}_m)$ with it's cell specific variant, $\bar{\Delta}_{mg} \log(\pi^{cc}_{mg})$. I also estimate models that account for cell level demand shocks.

I estimate models of the form

$$\Delta(emp)_{mg} = \sum \alpha^g + \beta\bar{\Delta}_{mg}\log(\pi^{cc}_{mg}) + f(emp^{t_0}_{mg})\gamma + X_m\delta + \epsilon_{mg} \tag{2.2}$$

where $\bar{\Delta}_{mg}\log(\pi^{cc}_{mg})$ and $\bar{\Delta}_{mg}\log(\pi^{cc}_{mg})$ are cell-specific measures of job suburbanization and demand shocks discussed in section 2.3.1.

Estimates are presented in the Appendix in Table 2.A.2. Notably, the cell level weighting appears to make little difference for the estimates. Moreover, the inclusion of cell specific demand shocks also has no noticeable impact on the results. This is true for weightings both on based industry and occupation. This is not surprising given that job suburbanization and growth are uncorrelated at the MSA and MSA by sector level.

2.4.3 Endogenous Migration

I argue above that a panel analysis has significant advantages over a cross-sectional analysis in testing predictions of the spatial mismatch hypothesis. However, the main analysis I have used a synthetic panel rather than a true panel of individuals. For the same reason that residential sorting is a concern for any cross-sectional analysis, the endogenous migration of households to and from MSAs may introduce a compositional bias in the synthetic panel analysis. The productivity of migrants to and from an MSA may vary systematically with job suburbanization, so that the correlation between job suburbanization and employment growth may in part reflect the changing composition of cells rather than within-person changes. For example, productive black households may move out of an MSA following widespread suburbanization, leaving less productive black households behind. While this type of endogenous migration would suggest that the spatial distribution of work is relevant for labor market outcomes, it complicates the interpretation of the coefficient estimates.

In assessing the role of migration, I again focus on the 1970-1980 period. I utilize the fact that the 1980 Census identifies an individual's place of residence and employment status in 1975 for half the sample. Hence, I observe a true panel of a large subsample of individuals over this 5-year period. I estimate models analogous to (2.1), except over a 5-year period, and compare the results when using a synthetic panel to results using a true panel. I construct the synthetic panel by assigning individuals to MSAs separately by where they live in 1975 and where they live in 1980. To form the true panel, I assign individuals to MSAs based only on where they live in 1975.

Unfortunately, I cannot observe spatial or sectoral distribution of work in 1975. Instead, I use the same measures of job suburbanization and job growth over the 10-year period. As a result, we should expect smaller coefficients here compared to 10-year analysis. In addition, even under random migration, the coefficients in the true panel analysis will be relatively attenuated for the same reason that attrition attenuates coefficient estimates. Households that move are mechanically less 'exposed' to job suburbanization in their assigned MSA than households that do not move. Under random migration, given that about 15% of individuals change MSAs over this period, and assuming coefficients of zero for moving households, this attrition would predict that the coefficients

Table 2.5: Job Suburbanization, Employment, and Migration, 1975-1980

	Δ log (Emp. Rate) True (1)	Synthetic (2)	Fraction Emigrate (3)	$\hat{\Delta}$ log (Emp. Rate) Baseline (4)
$\Delta \log(\pi^{cc})$	0.030	0.004	0.055*	0.008
	(0.030)	(0.039)	(0.024)	(0.005)
$\Delta \log(\pi^{cc}) \times$ black	0.147**	0.180*	-0.032	0.001
	(0.055)	(0.074)	(0.025)	(0.004)
Group FEs	Yes	Yes	Yes	Yes
Quadratic in Baseline	Yes	Yes	Yes	Yes
N Cells	2015	2015	2015	2015
N MSAs	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

in the true panel analysis will be about 0.85 times as large as the coefficients in the synthetic panel analysis.

I present estimates from the 5-year analysis in Table 2.5. In columns (1) and (2) the outcome is Δ log (Fraction Employed). In columns (3) and (4) the outcome is Δ (Fraction Employed). In columns (1) and (3) the analysis is conducted using a true panel; columns (2) and (4) are estimated using a synthetic panel. In general, the coefficients do not differ substantially. As would roughly be predicted by random migration, the coefficient on $\Delta \log(\pi_m^{cc}) \times black$ is about 15% smaller in magnitude under the true panel, and statistically significant using either the true or synthetic panel.

In addition, I test directly for selective migration by examining whether job suburbanization predicts migration flows and the composition of migrants. For migration flows I look at the fraction of a cell that emigrates out of the original MSA. To examine the composition of emigrants, I look at how the 1975 employment rate of a cell would change if the eventual migrants were excluded (I label this as $\Delta^M(emp)$). Columns (5)-(7) present coefficient estimates from this exercise. While job suburbanization predicts a small, marginally significant decrease in emigration for white cells, it is uncorrelated with changes in the composition of cells due to emigration.

2.4.4 Job versus Residential Suburbanization

Another concern with the above results is about interpretation. While to this point I have emphasized job suburbanization as the explanatory variable of interest, this measure may be closely related to residential suburbanization, which may have a direct effect on black relative employment for reasons discussed above that are orthogonal to spatial mismatch. Though it is difficult to imagine an exercise that can perfectly disentangle the effects of job and residential suburbanization, I

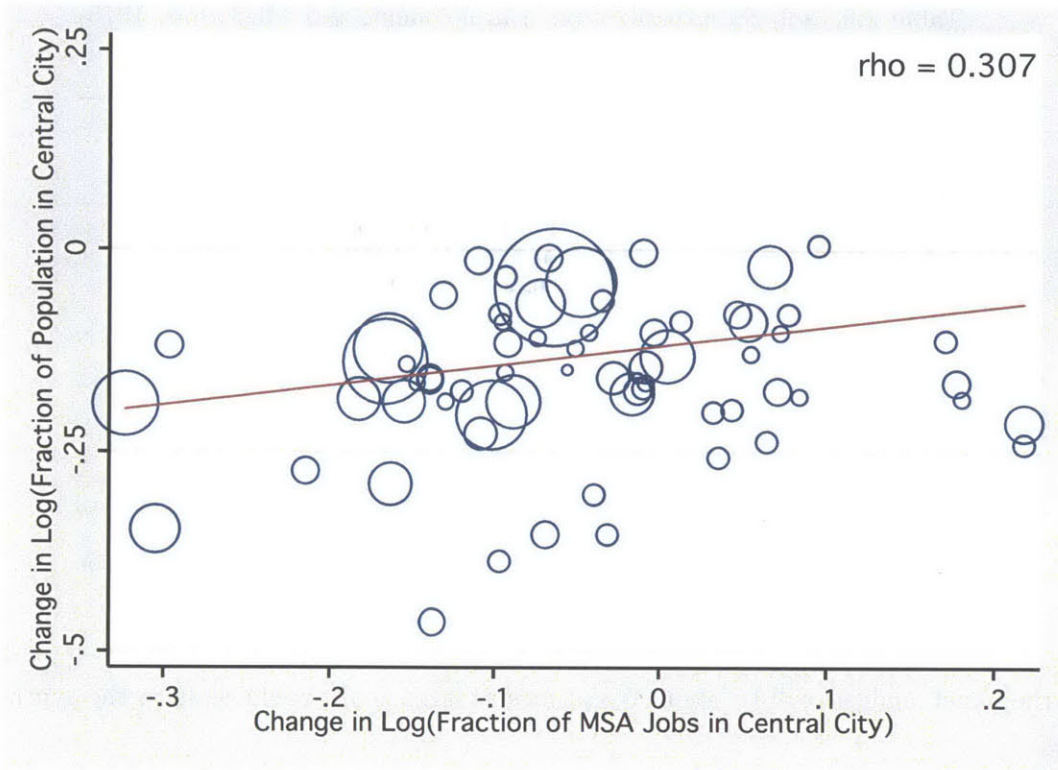

Figure 2.2: Job versus Residential Centralization, 1970-1980

Notes: This figure plots changes in job centralization against residential job centralization across MSAs from 1970 to 1980.

attempt to assess the potential role for residential segregation in generating the above results.[12]

I begin by describing the relationship between job and residential suburbanization across cities, measuring household centralization in a manner analogous to $\Delta \log \pi^{cc}$. I plot the relationship between MSA job and residential centralization from 1970 to 1980 in Figure 2.2 . Surprisingly, the relationship is not particularly strong: the weighted correlation coefficient is 0.31. While every MSA except Jackson, MS, undergoes residential suburbanization over this period, in many MSAs the fraction of work located in the central city increases.

To assess whether the above results are driven by the suburbanization of households rather than work *per se*, I estimate a variant of (2.1) that includes household centralization as a covariate. Results are presented in Table 2.6. All models include a quadratic in baseline employment. The results show that the relationship between residential suburbanization and black relative employment is quite different from the relationship between job suburbanization and black relative employment. In particular, while residential suburbanization predicts overall declines in employment rates, it has no detectable relationship with black relative employment. Conversely, conditional on residential suburbanization, there remains a strong negative relationship between job suburbanization and black

[12]Note that the population studied here is older than school-aged in 1970 and so their human capital should be unaffected by changes in school quality over this period.

Table 2.6: Job versus Residential Suburbanization, 1970-1980

	Δ log (Emp. Rate)		Δ log (Earnings)	
	(1)	(2)	(3)	(4)
$\Delta \log(\pi^{cc}_{res})$	0.160*	0.178**	0.038	0.061
	(0.063)	(0.060)	(0.098)	(0.097)
$\Delta \log(\pi^{cc}_{res}) \times$ black	0.010	-0.117	-0.086	-0.160
	(0.109)	(0.104)	(0.126)	(0.118)
$\Delta \log(\pi^{cc}_{job})$		-0.058		-0.066
		(0.048)		(0.065)
$\Delta \log(\pi^{cc}_{job}) \times$ black		0.295**		0.197*
		(0.067)		(0.092)
Group FEs	Yes	Yes	Yes	Yes
Quadratic in Baseline	Yes	Yes	Yes	Yes
N Cells	1944	1944	1944	1944
N MSAs	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

relative employment.

2.4.5 Highways as Exogenous Variation

I argue in section 2.2 that there are three factors that can induce job suburbanization: (1) declines in the relative non-labor costs of operating in the suburbs; (2) the suburbanization of labor; (3) shocks to labor productivity so that labor that remains in the central city becomes relatively less productive. Some combination of these general factors produced observed variation in job suburbanization over this period. However, for the observed negative relationship between job suburbanization and black relative employment to be taken as evidence in favor of the spatial mismatch hypothesis, it is important to distinguish between these underlying factors. In particular, unobserved shocks to the productivity of black labor (driven by supply or demand factors) could induce firms to migrate *and* produce black employment declines, even in the absence of a spatial mismatch mechanism. For example, the emergence of crack cocaine markets in the 1980s and 1990s could potentially explain both the deterioration of some black communities and the relocation of employers from the central city (Evans et al. 2012).

To test whether the observed relationship between job suburbanization and declines in black relative employment is driven entirely by such productivity shocks, I identify a source of variation in job suburbanization that is plausibly orthogonal to such shocks. The ideal source of variation would only produce job suburbanization directly through the non-labor entry and operating costs channel (and potentially generating employee suburbanization indirectly), rather than by inducing employee suburbanization directly. This is because, as in the Coulson et al. framework, job suburbanization induced by employee suburbanization may not have any effect on black employment, regardless

of whether spatial mismatch is relevant. In addition, the ideal source of variation would satisfy a condition similar to the instrumental variables exclusion restriction: it would only (potentially) affect black employment by changing the spatial distribution of work.

In practice, these two criteria are difficult to satisfy. Instead, I exploit a previously used source of variation in suburbanization: the interstate highway system (Baum-Snow 2007, Michaels 2008). Highways are an imperfect source of variation for two reasons. First, while highways reduce the operating costs firms faced in the suburbs, but also make suburban residence more attractive for individuals directly. Second, highways may affect black employment through channels other than changing the spatial distribution of work, for example, by reducing commuting costs. For these reasons, I will *not* argue that highways serve as a valid *instrument* for job suburbanization in this context. Instead, I focus on the reduced form relationship between highways and employment rates. I will argue that highways predict job suburbanization yet are orthogonal to supply and demand shocks particular to black labor. Hence, if highways predict black relative employment declines, the observed negative relationship between job suburbanization and black relative employment *must not be driven entirely by unobserved productivity shocks.*

Baum-Snow (2007) documents that the U.S. interstate highway system played an important role in post-war residential suburbanization. With nearly all construction occurring between 1956 and 1980, the interstate highway system would ultimately span over 40,000 miles. The highway system was originally designed to connect major metropolitan areas, serve U.S. national defense, and connect major routes in Canada and Mexico. Using plausibly exogenous variation in highway construction across MSAs, Baum-Snow finds that one new highway passing through a central city reduces its population from 1950 to 1990 by about 18%. These effects are substantial: they imply that the interstate highway system accounts for about 1/3 of the decline in central city population relative to total MSA growth over this period.

Highways can potentially increase suburbanization through several mechanisms. First, they decrease transportation costs for both firms and households. For firms, highways make physical proximity to other transportation hubs (e.g. ports and rail stations) and upstream or downstream firms less important, allowing them to take advantage of cheaper land and other suburban amenities. Michaels (2008) argues that, as highway construction was nearing completion, trucks became the primary mode of shipping goods within the United States. For households, highways reduce the costs of commuting to central work and access to other central city amenities from a suburban residence. These direct effects on transportation costs may also have feedback effects. By increasing the number of firms and households in the suburbs, they make these areas more attractive for other firms and households. In particular, firms may follow suburbanizing workers and achieve agglomeration economies in the suburbs; households may follow firms and take advantage of improved Tiebout sorting facilitated by highways.

In addition, highways have other important effects on local economies. Michaels (2008) shows that highways increased trade for exposed rural counties, and raised the relative demand for skilled manufacturing workers in skill-abundant counties while reducing it elsewhere. Duranton et al.

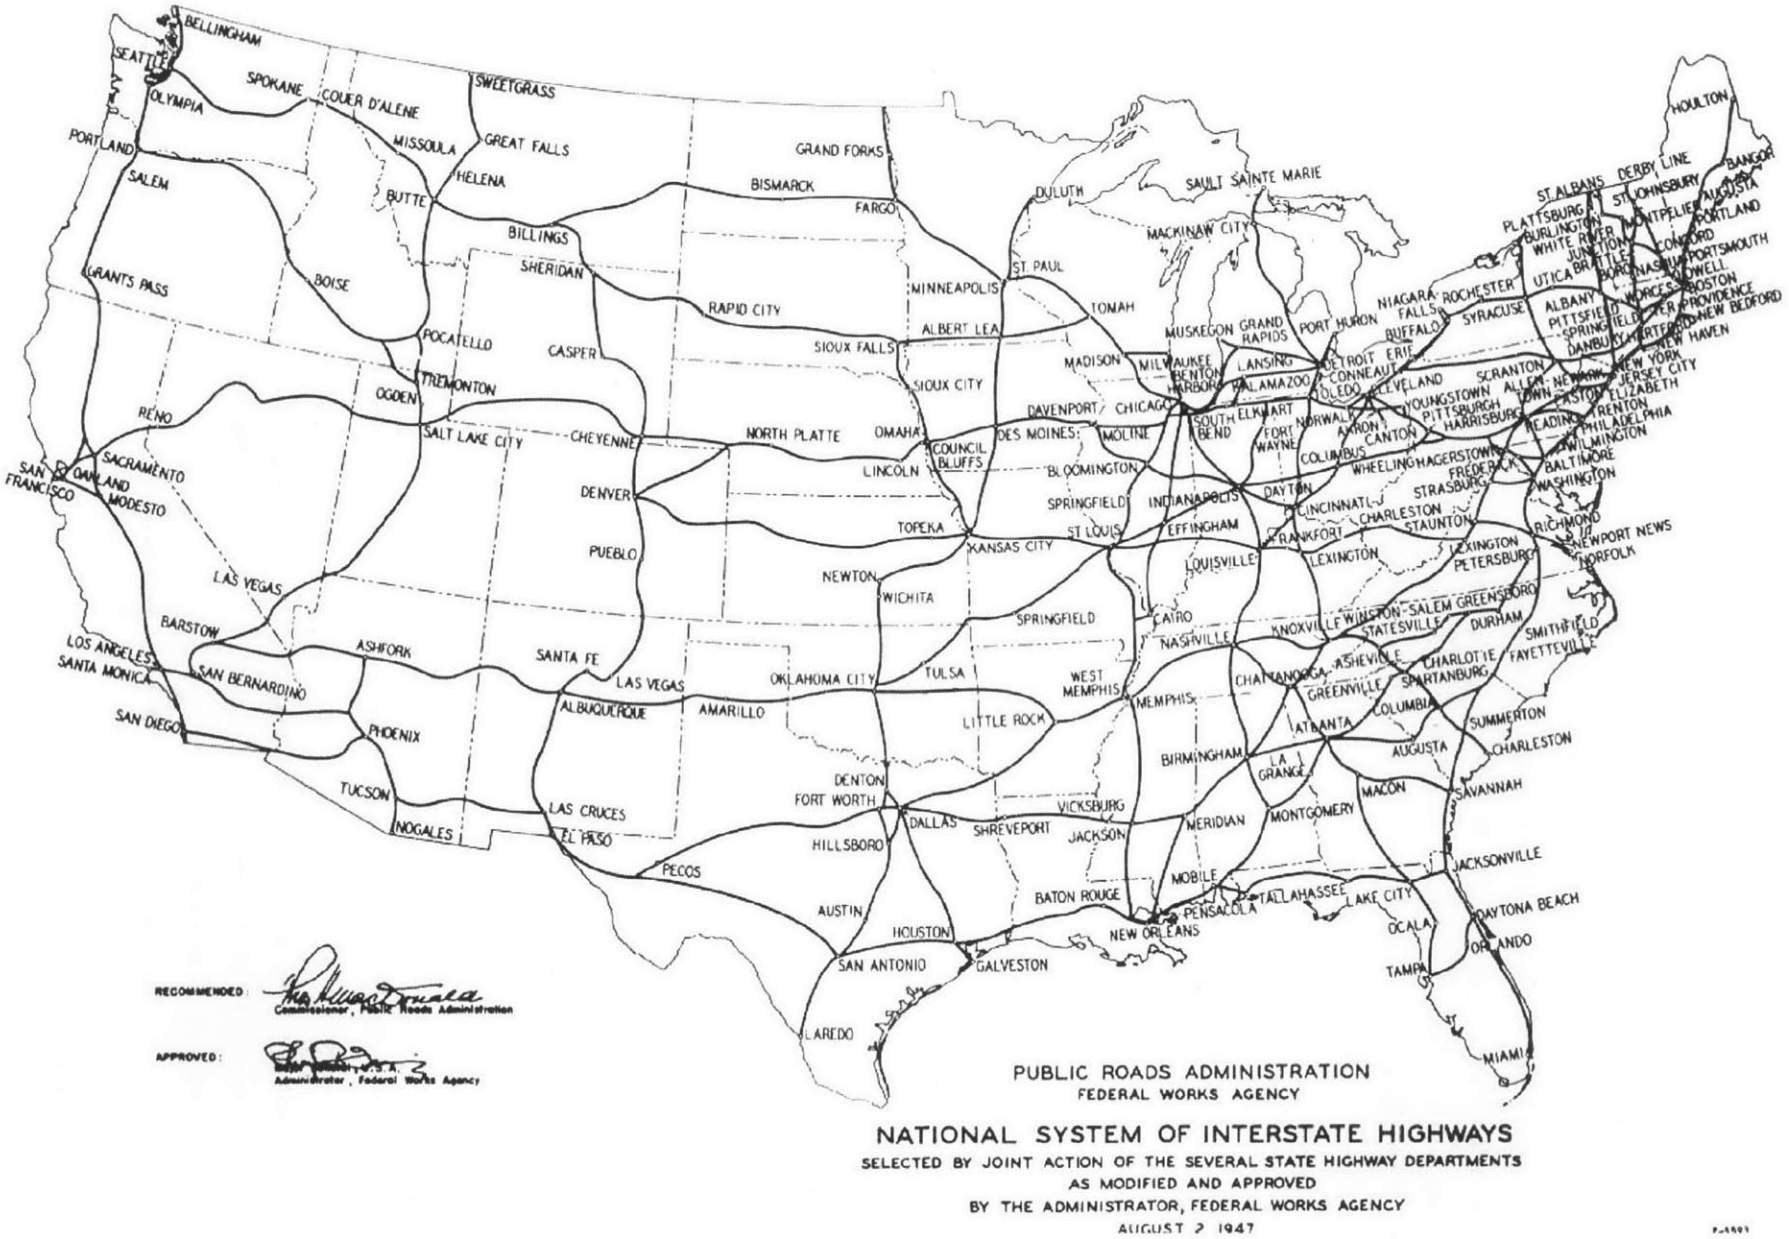

Figure 2.3: 1947 Interstate Highway Plan

(2011) find that highways increase the weight of a city's exports and that cities with more highways specialize in sectors producing heavy goods. Duranton and Turner (2011) find that interstate highways increase MSA growth from 1983 to 2003.

Because highways have such varied effects, they are unlikely to serve as a valid instrument for job suburbanization in this context. In particular, they may affect black relative employment through channels other than the spatial distribution of work. Still, highway assignment, planned in the 1940's and 50's, should be orthogonal to 1970's labor supply shocks that disproportionately affect black labor. I also show that they appear to be unrelated to demand shifters that would particularly affect black labor over this period.[13]

A potential concern with this approach is that highway assignment may be determined endogenously. As Baum-Snow (2007) notes, the interstate highway system was likely in part designed to facilitate local commuting and local economic development in particular regions. The highway system might have also been designed accounting for productivity shocks to black workers in the 1970's, though this seems less plausible. Actual highway construction across MSAs may be related to counterfactual MSA-level outcomes. To deal with this, I instrument for realized highway construction using the 1947 interstate highway plan as in Baum-Snow (2007) and other work.

In 1937, the Franklin D. Roosevelt administration began to plan an interstate highway system. In their recommended highway plan, the National Interregional Highway committee considered the distribution of population, manufacturing activity, agricultural production, the location of post-World War II employment, a strategic highway network drawn up by the War Department in 1941, the location of military and naval establishments, and interregional traffic demand, in that order. This led to the Federal Highway Act of 1944, which instructed the roads commissioner to develop an initial plan for a national interstate highway system. As specified by the legislation, the highways in the planned system were to be "... so located as to connect by routes as direct as practicable, the principal metropolitan areas, cities, and industrial centers, to serve the national defense, and to connect at suitable border points with routes of continental importance in the Dominion of Canada and the Republic of Mexico..." (as cited in Baum-Snow 2007). Importantly, the legislation makes no mention of local commuting or local economic development.

Major funding for the interstate highway system began with the Federal Highway Acts of 1956. The 1956 Interstate Highway Act expanded the 1947 plan and committed the federal government to pay 90 percent of the cost of construction. In particular, the 1956 plan incorporated additional highways that were designed for local purposes like commuting and development. Therefore, in some specifications I instrument for actual highway rays using highway rays planned in 1947. The first stage t-statistic is in excess of 5.

To begin, I explore the relationship between highways and several key outcomes: job suburbanization, job growth, and residential suburbanization. In all subsequent analyses, my measure of highway exposure is the number of interstate highway rays emanating from the central city in 1970.

[13] Baum-Snow (2007) finds that MSA non-white population and percent high school graduate cannot predict either observed highway construction or rays planned in 1947.

Table 2.7: Highways and Job Suburbanization, 1970-2000

	1970-1980		1970-1990		1970-2000	
	OLS	IV	OLS	IV	OLS	IV
Outcome: $\Delta \log(\pi^{cc})$	(1)	(2)	(3)	(4)	(5)	(6)
Rays	-0.036**	-0.038*	-0.057**	-0.076**	-0.076**	-0.110**
	(0.010)	(0.016)	(0.014	(0.022)	(0.020)	(0.028)
Rays × black	-0.005	-0.015~	-0.007	-0.025*	-0.010	-0.029~
	(0.006)	(0.009)	(0.007)	(0.013)	(0.010)	(0.017)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
CC Radius Control	Yes	Yes	Yes	Yes	Yes	Yes
N Cells	1870	1870	1278	1278	659	659
N MSAs	68	68	68	68	68	68

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

To determine the relationship between highways and job suburbanization and growth, I estimate specifications of the following form, where again coefficients can differ by race:

$$\Delta Y_{mg} = \gamma_1 rays_m^{1970} + \gamma_2 radius_m + f(emp^{1970})_{mg} + X_{mg}\gamma_3 + \sum I^g + \epsilon_{mg} \tag{2.3}$$

where $rays_m^{1970}$ denotes the number highway rays emanating from the central city of MSA m in 1970 and $radius_m$ is the radius of the central city, a key control in the analysis of Baum-Snow (2007). Intuitively, one must control for central city radius because it determines the extent to which sprawl is reflected in suburbanization measures. Again, some specifications include a vector X of additional MSA-level controls including regional indicators and the fraction of the MSA population that is black in 1970. Note that the average number of central city highway rays in 1970, weighted by the sample population size, is 3.9; the unweighted average is 3.[14]

Table 2.7 presents estimates where the outcome is $\Delta_m \log(\pi^{cc})$. I find that the stock of highways in 1970 predicts job suburbanization thereafter. In odd specifications I use actual interstate highway rays constructed as the explanatory variable of interest; in even specifications I instrument for highways constructed using highway rays included in the 1947 plan. In all columns, t_0 is 1970, while t_1 is 1980 in columns (1) and (2), 1990 in columns (3) and (4), and 2000 in columns (5) and (6). In the baseline OLS specification, 1 highway ray emanating from the central city is associated with a 3.6%, 5.7%, and 7.6% decrease in the fraction of MSA jobs located in the central city by 1980, 1990, and 2000. When I instrument for highways, the point estimates increase somewhat, particular when weighting by black population, though they are less precise.

In the Appendix in Table 2.A.3 the outcome is $\bar{\Delta}_{mg} \log(T)$, a cell-specific weighted average of MSA sector growth. Again, actual interstate highway rays constructed is the explanatory variable in odd columns, while highway rays are instrumented in even columns. When using industry weights, there is little relationship between highways and changes in relative demand. When using occupation

[14]Note that if a highway passes through a central city, this counts as 2 rays.

weights, a slight negative relationship emerges for black cells during the 1980's.

To determine the relationship between highways and residential suburbanization and segregation, I estimate MSA-level regressions rather than group-level regressions. I do this because I cannot observe whether individuals live in the central city in the 1970 Census. Instead, I use population data from the *County and City Data Books* (CCDB), which report decennial Census data aggregated to counties and cities of at least 25000 inhabitants. For the household suburbanization and segregation analysis, I estimate specifications of the following form:

$$\Delta Y_m = \delta_1 rays_m^{1970} + \delta_2 \Delta \log(pop)_m + \delta_3 radius_m^{cc} + \epsilon_m \tag{2.4}$$

where Y denotes either $\log \pi_{res}^{cc}$, where π_{res}^{cc} is the fraction of an MSA's population living in the central city, or the dissimilarity index and *pop* denotes total MSA population. I estimate this model for the whole population and separately by race when the outcome is household suburbanization.

I present estimates in Table 2.8. Highway rays predict the suburbanization of white households, but not black households. This disparity is consistent with a central premise of the spatial mismatch hypothesis: that black households faced significant additional barriers to suburban residence over this period. Each ray predicts a 5.4%, 9.0%, and 11.3% decline in the central city share of the white population by 1980, 1990, and 2000. By contrast, highways predict only a negligible decline in the black population. The effects are larger when I instrument for highway construction, though the effect on the black central city population remains statistically insignificant. Highway rays increase subsequent, with each ray predicting an increase of 0.010 to 0.016 in a city's dissimilarity index.

Finally, I estimate the reduced form effect of highways on black relative employment and earnings.[15] I estimate specifications of the same form as (2.3), where Y is now log employment rate or log earnings. I present estimates in Table 2.9.

Across specifications, highways predict black relative employment declines. Odd specifications are OLS models, while I instrument for highways in even specifications. All specifications include a quadratic in baseline employment. The pattern of coefficients is consistent across specifications and outcomes. While highways predict increases in the employment rates and earnings of whites, they predict relative decreases in these labor market outcomes for blacks. For the employment outcomes, the magnitudes of the coefficients are stable across specifications. In the OLS models, each additional highway ray predicts a 1% increase in white employment rates and weeks worked, but about 2% relative decline in black employment and weeks worked. Coefficients from the IV models are similar. Strikingly, while a racial gap in employment rates emerges from 1970 to 1980, this remains roughly constant through 2000. As in section 2.4.1, the magnitude of the coefficients in the earnings models is sensitive whether I control for baseline earnings.

To assess the robustness of these estimates, I re-estimate this set of specifications including additional MSA-level controls for regional fixed effects and the initial black share of the MSA population as in section 2.4.1. The results are presented in the Appendix. Again, the general

[15] I also explore how the stock of highway rays in 1970 relates to other baseline group and MSA characteristics, analogous to the exercise in section 2.4.1. The results are presented in the Appendix.

Table 2.8: Highways, Residential Suburbanization, and Segregation, 1970-2000

	1970-1980		1970-1990		1970-2000	
	OLS	IV	OLS	IV	OLS	IV
Outcome: $\Delta \log(\pi^{cc}_{res})$, White	(1)	(2)	(3)	(4)	(5)	(6)
Rays	-0.054**	-0.077**	-0.090**	-0.135**	-0.113**	-0.173**
	(0.013)	(0.021)	(0.024)	(0.038)	(0.034)	(0.056)
$\Delta \log(pop)$	-0.290**	-0.358**	-0.252*	-0.327**	-0.184	-0.248*
	(0.087)	(0.075)	(0.117)	(0.099)	(0.124)	(0.110)
Outcome: $\Delta \log(\pi^{cc}_{res})$, Black	(7)	(8)	(9)	(10)	(11)	(12)
Rays	0.003	-0.007	-0.003	-0.031	-0.019	-0.051
	(0.010)	(0.017)	(0.021)	(0.029)	(0.032)	(0.044)
$\Delta \log(pop)$	-0.143	-0.166~	-0.225**	-0.249**	-0.305**	-0.319**
	(0.099)	(0.087)	(0.078)	(0.077)	(0.103)	(0.102)
CC Radius Control	Yes	Yes	Yes	Yes	Yes	Yes
N MSAs	68	68	68	68	68	68
Outcome: $\Delta \log(Dissimilarity)$	(13)	(14)	(15)	(16)	(17)	(18)
Rays	0.010	0.023~	0.042**	0.044*	0.049**	0.051*
	(0.013)	(0.012)	(0.011)	(0.020)	(0.015)	(0.024)
CC Radius Control	Yes	Yes	Yes	Yes	Yes	Yes
N MSAs	65	65	65	65	65	65

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size. Data from County and City Data Books.

Table 2.9: Highways and Labor Market Outcomes, 1970-2000

	1970-1980		1970-1990		1970-2000	
	OLS	IV	OLS	IV	OLS	IV
Outcome: $\Delta \log$ (Employment Rate)	(1)	(2)	(3)	(4)	(5)	(6)
Rays	0.009*	0.011*	0.010**	0.014*	0.014**	0.018*
	(0.003)	(0.005)	(0.003)	(0.006)	(0.005)	(0.007)
Rays	-0.022**	-0.021*	-0.018**	-0.018~	-0.019**	-0.026**
× black	(0.005)	(0.009)	(0.007)	(0.010)	(0.006)	(0.010)
Outcome: $\Delta \log$ (Earnings)	(7)	(8)	(9)	(10)	(11)	(12)
Rays	0.009~	0.019**	0.011	0.007	0.022*	0.017
	(0.005)	(0.007)	(0.007)	(0.012)	(0.010)	(0.013)
Rays	-0.019**	-0.021*	-0.020~	-0.017	-0.017	-0.020
× black	(0.006)	(0.010)	(0.010)	(0.014)	(0.010)	(0.017)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic in Baseline	Yes	Yes	Yes	Yes	Yes	Yes
CC Radius Control	Yes	Yes	Yes	Yes	Yes	Yes
N Cells	1870	1870	1278	1278	659	659
N MSAs	68	68	68	68	68	68

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

results are generally unchanged; the coefficient estimates for $rays \times black$ are somewhat muted and estimated less precisely, though still statistically significant, but the relationship between highway rays and job suburbanization is also weaker.

I also repeat the 5-year panel analysis as in Section 2.4.3 using variation in highways. I report the analogous IV specifications in the Appendix. As before, the relevant point estimates are smaller in the true panel but still statistically significant at the 5% level. The proportional gap between the true and synthetic panel point estimates is larger here, but I cannot reject equality of the point estimates across the two specifications. Endogenous migration does not appear to be driving the relationship between highways and black relative employment declines.

2.5 How Do Jobs Suburbanize?

While residential suburbanization is typically thought to be driven by households and individuals moving out of the central city and into the suburbs, establishment relocations are less common. Instead, the spatial distribution of work mostly shifts with the entry and exit of firms; the fraction of work in the suburbs can increase if the rate of establishment exit is higher in the central city, the rate of entry is higher in the suburbs, or some combination of the two. To this point I have focused on the net effect; however, whether job suburbanization occurs primarily through relative exit or relative entry affects the interpretation of the relationship between job suburbanization and black relative employment. In particular, 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 correlation may simply reflect the effects of job displacement rather than the spatial distribution of work *per se*. In this section I decompose suburbanization into the relative exit and relative entry channels, and investigate the extent that baseline results may be driven by the effects of job displacement.

To undertake this analysis, I use unique EEO-1 form establishment-level panel data from 1971-1981, collected by the U.S. Equal Employment Opportunity Commission (EEOC). As part of the Civil Rights Act of 1964, firms meeting certain size requirements are required to complete EEO-1 forms annually and submit them to the EEOC. Firms are to report their overall racial and gender composition, the racial and gender composition of each of their establishments meeting certain size requirements, as well as the industry, size, and address of those establishments. Firms are also required to report racial and gender composition separately by 9 major occupation groups. Over this period, all firms with 50 or more employees were required to submit EEO-1 forms, and were required to file a separate report for each establishment with at least 50 employees. For federal contractors, both cutoffs were 25 employees. Establishments are consistently identified with firm and establishment identifiers.

I present descriptive statistics for the EEO-1 data in Appendix. Due to the size requirements, establishments in the EEO-1 data are not representative of all U.S. establishments.[16] In 1980, the

[16]In addition, some firms fail to submit required EEO-1 forms.

data cover about 92000 establishments and nearly 16 million employees in MSAs I include in my analysis. Industries that tend to have large establishments, e.g. manufacturing, are overrepresented, while industries that tend to have small establishments, e.g. services, are underrepresented. In general, the entry and exit behavior of medium and large firms differ from that of small firms. There is some evidence that larger firms are more likely to hire black workers (Holzer 1998). However, the focus here is on variation in entry and exit rates across MSA central cities and suburbs rather than *absolute* entry and exit, so this lack of representativeness should be less of a concern. The (weighted) correlations between measures of $\Delta \log(\pi^{cc})$ and $\Delta \log(T)$ in the EEOC and Census data are 0.61 and 0.77.

I decompose job growth in MSA central cites and suburbs into the establishment exit rate, entry rate, and growth rate of incumbents. Let N_t denote the initial number of jobs in year t, X_t denote the initial number of jobs in establishments that exit during year t, let E_t denote the final number of jobs in establishments that enter during year t, and let g_t^i denote the growth rate of incumbent firms in year t. Then

$$\begin{aligned} N_{t+1} &= N_t - X_t + E_t + g_t^i(N_t - X_t) \\ \Rightarrow \frac{N_{t+1} - N_t}{N_t} &= -\frac{X_t}{N_t} + \frac{E_t}{N_t} + g_t^i\left(1 - \frac{X_t}{N_t}\right) \\ \Rightarrow g_t &= -x_t + e_t + g_t^i(1 - x_t) \\ &\approx -x_t + e_t + g_t^i \end{aligned}$$

where g is the overall annual growth rate, x is the exit rate, e is the entry rate, and g^i is the incumbent growth rate.

Rather than directly compute these rates at an annual frequency and take a geometric average, I use the annual rates implied by the 10-year rates. I do this because the EEO-1 data are somewhat noisy. Some establishments change locations and then change back in a manner that appears unlikely to reflect real moves. Some of these instances, though not all, appear to be firms mislabeling their establishments. Some establishments exit and re-enter the data set. While this can occur mechanically if the size of the establishment is near the threshold or the federal contractor status of the parent firm is changing, these explanations do not appear to be relevant in many instances. To deal with this, I focus on establishments that are less likely to be recorded with error, and this is easier to do at a lower frequency. In particular, I identify an establishment as present in 1971 in a given MSA's central city or suburbs if it is located in that area in 1971 and the next year it is recorded in the data (establishments that are only present for one year from 1971-1981 are excluded). Similarly, I identify an establishment as present in 1981 in a given MSA's central city or suburbs if it is located in that area for two consecutive years from 1979 to 1981.

To compute the annual overall growth, exit, and incumbent growth implied by their 10-year counterparts, I calculate the constant annual rate consistent with overall growth, and the exit and growth of establishments identified as present in 1971. To calculate overall growth, I find g such

Table 2.1: Job Growth in the EEO-1 Data, 1971-1981

	Exit Rate (x)	Entry Rate (e)	Incumbent Growth Rate (g^i)	Overall Growth Rate (g)
Central City	0.083	0.069	0.003	-0.012
	(0.018)	(0.013)	(0.011)	(0.026)
Suburbs	0.074	0.088	0.002	0.017
	(0.024)	(0.022)	(0.018)	(0.028)
Within MSA Difference	0.012	-0.025	-0.002	-0.039
	(0.029)	(0.023)	(0.017)	(0.031)
N MSAs	64	64	64	64

that $N_{1971}(1+g)^{10} = N_{1981}$. I calculate incumbent growth analogously. For calculating the exit rate, let I_t^{1971} denote the number of jobs in 1971 among incumbent firms (present in 1971) that are still around in year $t \geq 1971$. I compute the exit rate as x, where x satisfies $I_{1971}^{1971}(1-x)^{10} = I_{1981}^{1971}$. With values for g, x, and g^i, I back out the entry rate, e.

I present summary results from decomposing job suburbanization in Table 2.1. I show the mean and standard deviation exit rates, entry rates, incumbent growth rates, and overall growth rates separately for central cities and suburbs. These statistics are weighted by the number of workers in the area (i.e. central city or suburb) in 1971. I also show the average *within*-MSA difference in these rates (central city minus suburbs), weighting by the number of workers in the MSA in 1971. On average, suburbs have lower exit rates, higher entry rates, and comparable incumbent growth rates. Suburbs grew on average by 1.7% per year; central cities shrank by 1.2%. However, these averages mask differences across MSAs in these rates and the relative sizes of central cities and suburbs within MSAs. The within-MSA differences are more relevant for this analysis. Within MSAs, the gap in the central city and suburban growth rate is even larger at 3.9%. Across exit, entry, and incumbent growth rates, the largest differences are in entry rates. About 65% of the central city-suburb growth gap can be attributed to differential rates of establishment entry. In contrast, the exit rate differential is half as large.

However, though the central city-suburb gap is largest for the entry rate, MSA job suburbanization may still be significantly correlated with the central city exit rate or more importantly, the differential in. To check this, I regress the central city and suburb exit rate, entry rate, and growth rates on MSA job suburbanization and total MSA job growth. I also relate job suburbanization and total MSA job growth to the establishment exit rate faced by black and white workers. These are the exit rates of incumbent firms weighted by the number of black or white employees initially working at the establishment. Coefficient estimates are presented in Table 2.2.

Surprisingly, there is little relationship between MSA job suburbanization and central city and suburban entry rates. Suburbanization is negatively correlated to central city incumbent growth but uncorrelated to suburban incumbent growth. Most importantly, job suburbanization is positively correlated with central city exit rates and negatively correlated with suburban exit rates. This is also true of both the black and white central city and suburban exit rates. However, these effects are offsetting so that job suburbanization does *not* predict black or white MSA-wide exit

Table 2.2: Decomposing Job Suburbanization, 1971-1981

	Exit Rate (x)	Entry Rate (e)	Incumbent Growth Rate (g')	Overall Growth Rate (g)	Black Exit Rate	White Exit Rate
Central City						
$\Delta \log(\pi^{cc})$	-0.048**	0.015	0.036**	0.099**	-0.047**	-0.052**
	(0.013)	(0.010)	(0.008)	(0.000)	(0.016)	(0.013)
$\Delta \log(T)$	-0.045**	0.030**	0.025**	0.099**	-0.038**	-0.046**
	(0.009)	(0.007)	(0.006)	(0.000)	(0.011)	(0.009)
Suburbs						
$\Delta \log(\pi^{cc})$	0.086**	-0.008	0.006	-0.088**	0.115**	0.083**
	(0.018)	(0.019)	(0.015)	(0.016)	(0.027)	(0.018)
$\Delta \log(T)$	-0.029*	0.057**	0.043**	0.128**	-0.022	-0.030*
	(0.013)	(0.014)	(0.011)	(0.012)	(0.019)	(0.013)
Total MSA						
$\Delta \log(\pi^{cc})$					-0.002	0.004
					(0.016)	(0.013)
$\Delta \log(T)$					-0.032**	-0.035**
					(0.011)	(0.009)
N MSAs	64	64	64	64	64	64

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses. Regression weighted by area's total EEO-1 employment in 1971.

rates. Given that I pool workers by MSA in the above analysis (rather than at the MSA by central city status level), this suggests that the results in section 2.4.1 are unlikely to be driven by job displacements. Moreover, this provides additional evidence that suburbanization is unlikely to be driven by unobserved changes to black relative productivity.

I also repeat this analysis using highways as the explanatory variable, controlling for central city radius. The results are presented in the Appendix. Again, while highways do predict central city exit rates, they do not predict MSA-wide black or white exit rates, suggesting that job displacements are not driving the results in section 4.4.

2.6 Conclusion

For several decades, spatial mismatch has been a commonly cited cause of persistently high black unemployment. However, previous tests of spatial mismatch are plagued by the endogeneity of household and firm location. I account for these endogeneity problems using panel methods and exogenous variation in highway construction. My estimates imply that from 1970 to 2000, for every 10% decrease in the fraction of MSA jobs located in the central city, black relative employment rates declined by 1.4-2.1%, while relative earnings decline by 1.1-2.3%. These estimates are not driven by selective migration or residential suburbanization. They are also not driven by unobserved shocks to labor productivity: 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. The negative relationship between job suburbanization and black relative employment emerges in the 1970's, but stagnates thereafter. The results suggest that job suburbanization was an important determinant of black labor market outcomes over the 1970's and its initial impact persisted.

To evaluate the importance of job suburbanization in producing labor market inequality over this period, I use the estimated coefficients to examine the overall contribution of job suburbanization to realized changes in relative employment rates. From 1970 to 1980, the proportion of black men living in the MSAs analyzed here that were active in the labor market decreased from 0.870 to 0.754. For white men, the proportion decreased from 0.947 to 0.868. In contrast, the share of women active in the labor market increased: for black women, from 0.574 to 0.601; for white women from 0.480 to 0.553. Altogether, the share of black adults active in the labor market decreased from 0.704 to 0.668, while that share of white adults remained constant at about 0.706. Hence, the black active share decreased by 5% or 3.6 percentage points over this period, while the white active share was unchanged. Given that the fraction of MSA jobs located in the central city declined by 9.5% for the average black adult, a β estimate of 0.262 implies that job suburbanization decreased the black active share by about 2.5% relative to the white active share over this period, or 50% of the overall decline.

2.A Appendix

Table 2.A.1: Robustness Checks: Job Suburbanization and Labor Market Outcomes, 1970-1980

	1970-1980				1970-2000			
	Δ (Emp. Rate)	Δ log (Emp. Rate)			Δ (Emp. Rate)	Δ log (Emp. Rate)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \log(\pi^{cc})$	-0.014	-0.006	0.006	0.017	-0.013	-0.029	-0.030	-0.030
	(0.029)	(0.048)	(0.055)	(0.056)	(0.019)	(0.029)	(0.029)	(0.030)
$\Delta \log(\pi^{cc}) \times$ black	0.163**	0.193**	0.165**	0.187**	0.063**	0.097*	0.098*	0.114*
	(0.038)	(0.061)	(0.059)	(0.063)	(0.019)	(0.042)	(0.040)	(0.045)
Fraction Black			0.120	0.061			0.120	0.054
			(0.115)	(0.104)			(0.146)	(0.131)
Fraction Black × black			-0.342*	-0.250~			-0.262	-0.174
			(0.134)	(0.130)			(0.195)	(0.202)

	1970-1980				1970-2000			
	Δ (Earnings)	Δ log (Earnings)			Δ (Earnings)	Δ log (Earnings)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \log(\pi^{cc})$	-801	-0.002	0.015	-0.001	2132	-0.028	-0.032	-0.026
	(685)	(0.067)	(0.073)	(0.073)	(1941)	(0.041)	(0.037)	(0.037)
$\Delta \log(\pi^{cc}) \times$ black	1187	0.139	0.120	0.153~	192	0.078	0.083	0.088
	(737)	(0.089)	(0.092)	(0.086)	(1198)	(0.071)	(0.066)	(0.064)
Fraction Black			0.213	0.292*			0.482*	0.422~
			(0.129)	(0.136)			(0.235)	(0.230)
Fraction Black × black			-0.250	-0.418*			-0.570~	-0.550~
			(0.185)	(0.175)			(0.310)	(0.320)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Quadratic in Baseline	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N Cells	1944	1944	1944	1944	686	686	686	686
N MSAs	71	71	71	71	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level. Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

Table 2.A.2: Job Suburbanization, Sector Composition, Labor Market Outcomes, 1970-1980

	Δ log (Emp. Rate) Industry Weighted			Δ log (Emp. Rate) Occupation Weighted		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log(\pi^{cc})$	-0.044	-0.027	-0.024	-0.036	-0.022	-0.015
	(0.046)	(0.053)	(0.054)	(0.043)	(0.049)	(0.052)
$\Delta \log(\pi^{cc}) \times$ black	0.226**	0.263**	0.263**	0.213**	0.236**	0.217**
	(0.059)	(0.070)	(0.072)	(0.059)	(0.068)	(0.065)
θ			-0.238			0.009
			(0.158)			(0.006)
$\theta \times$ black			0.395*			-0.002
			(0.186)			(0.008)

	Δ log (Earnings) Industry Weighted			Δ log (Earnings) Occupation Weighted		
	(7)	(8)	(9)	(10)	(11)	(12)
$\Delta \log(\pi^{cc})$	-0.022	-0.035	-0.029	-0.040	-0.056	-0.048
	(0.058)	(0.069)	(0.067)	(0.052)	(0.059)	(0.059)
$\Delta \log(\pi^{cc}) \times$ black	0.322**	0.134	0.130	0.314**	0.125~	0.130~
	(0.089)	(0.083)	(0.083)	(0.081)	(0.072)	(0.073)
θ			-0.393			0.002
			(0.319)			(0.007)
$\theta \times$ black			0.497*			-0.006
			(0.233)			(0.007)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
Quadratic in Baseline	No	Yes	Yes	No	Yes	Yes
N Cells	1944	1944	1944	1944	1944	1944
N MSAs	71	71	71	71	71	71

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

Table 2.A.3: Highways and Sector Growth, 1970-2000

	1970-1980		1970-1990		1970-2000	
Outcome: $\bar{\Delta}_{mg}\log(T)$	OLS	IV	OLS	IV	OLS	IV
Industry Weighted	(1)	(2)	(3)	(4)	(5)	(6)
Rays	-0.000	0.000	-0.001	-0.003	0.002	-0.003
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)
Rays × black	0.009~	0.005	0.004	-0.002	-0.000	-0.008
	(0.005)	(0.006)	(0.006)	(0.007)	(0.005)	(0.007)
Outcome: $\bar{\Delta}_{mg}\log(T)$						
Occupation Weighted	(7)	(8)	(9)	(10)	(11)	(12)
Rays	-0.002	-0.003	0.001	-0.001	0.003	-0.000
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.006)
Rays × black	-0.002	-0.002	-0.013**	-0.013*	-0.015**	-0.015~
	(0.004)	(0.006)	(0.003)	(0.006)	(0.004)	(0.008)
Group FEs	Yes	Yes	Yes	Yes	Yes	Yes
CC Radius Control	Yes	Yes	Yes	Yes	Yes	Yes
N Cells	1870	1870	1278	1278	659	659
N MSAs	68	68	68	68	68	68

Notes: * Denotes statistical significance at $p < 0.05$ level. ** Denotes statistical significance at $p < 0.01$ level.
Standard errors in parentheses, clustered at the MSA level. Regression weighted by cell size.

Chapter 3

Optimal Social Insurance with Heterogeneity

3.1 Introduction

In this paper we investigate 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.[1] This framework has been applied extensively in both empirical and theoretical work on social insurance (e.g., Gruber 1997; Chetty and Looney 2006; Chetty and Saez 2010; Landais et al. 2010; Gross and Notowidigdo 2011; Kroft and Notowidigdo 2011; Schmieder et al. 2012). One potential shortcoming of the Baily (1978) and Chetty (2006) 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.[2] 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 ex-

[1] 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.

[2] Chetty (2006) footnote 8, page 1894.

tend 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 unemployment insurance is more valuable if more risk averse agents are subject to larger risks. We refer to this as the *covariance effect.*[3] 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%.

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 unemployment insurance. Without knowing the joint distribution of risk preferences, ex-ante risk, and ability to smooth consumption, however, 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.

There is growing evidence of risk preference heterogeneity in insurance settings. In particular, a recent literature documents heterogeneity in risk preferences in insurance markets estimated using market outcomes. Cohen and Einav (2007) find evidence of substantial variation in absolute and relative risk aversion using observational data on deductible choice in Israeli auto insurance contracts. Barseghyan et al. (2011) and Einav et al. (2011) explore the stability of individual risk preferences across different insurance and investment domains and find evidence that person-specific risk preferences are significantly correlated across settings. Einav et al. (2011) also find that an individual's

[3]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).

revealed risk preferences in every other insurance domain predict that individual's choice in a given insurance domain better than demographic characteristics. Together, these findings suggest that there exists substantial heterogeneity in risk preferences, even within demographic subgroups. We show that this heterogeneity has substantive implications for the Baily-Chetty approach to welfare analysis.

Our study is closely related to Chetty and Saez (2010), who apply a sufficient statistic approach to characterize the welfare gains from social insurance when the private sector provides partial insurance. In this setting, the validity of the Baily-Chetty formula depends crucially on whether private insurance generates moral hazard. If so, the standard formula must be modified to account for fiscal externalities; otherwise, the formula is unaffected. In contrast, while we do not account for fiscal externalities our results highlight the importance of accounting for private insurance markets because they may affect the covariance of risks faced by workers and risk preferences.

This paper also relates to a new literature on tests for efficient risk sharing across households when household preferences for risk are heterogeneous. Standard tests are based on the idea that, under full risk sharing, household consumption should not respond to idiosyncratic shocks after accounting for aggregate shocks. Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) show that, under heterogeneity, such tests may reject efficiency even if households share risk efficiently. In particular, when some households are less risk averse than others, it is Pareto-efficient for those who are less risk averse to bear more aggregate consumption risk (Diamond 1967; Wilson 1968). Acknowledging this, both sets of authors derive tests of efficient risk sharing that allow for risk preference heterogeneity.[4] Similar to the present paper, this literature highlights that it is important to understand the relationship between risks faced by agents and risk preferences when preferences may be heterogeneous.

In the next section, we review the two-period model studied by Chetty (2008) and use this model to illustrate the basic sufficient statistics approach. In section 3.3, we propose a tractable approach to calculating aggregate welfare for heterogeneous agents. In section 3.4, using this approach we consider two separate cases: one where taxes are set to be actuarially fair within individual and the other where taxes are set uniformly, allowing for expected transfers between workers. We also discuss practical issues that arise in implementing our formula. In section 3.5, we show in a simple calibration exercise that the covariance effect is plausibly large. In section 3.6, we extend our results to cover benefits and taxes proportional to wages and higher order approximations to the utility function, and show that analogs of our baseline results hold in a heterogeneous version of the rich dynamic model studied by Chetty (2006). Section 3.7 concludes.

3.2 The Baily-Chetty Formula

We begin by reviewing the basic sufficient statistics approach of Baily (1978) and Chetty (2006) as applied to unemployment insurance. The starting point of our analysis is the 2-period job search

[4]Moreover, the authors are able to reject the null that models constraining household risk preferences to be homogeneous are correctly specified using data from the U.S. and India, respectively.

model described in Chetty (2008), which can be used to frame the analyses of Baily (1978) and Chetty (2006). Suppose an agent or worker becomes unemployed at $t = 0$ with assets A. The agent chooses search effort s, where s is normalized to be the probability that the agent finds a job. Let $\psi(s)$ denote the cost of search effort, where $\psi(\cdot)$ is strictly increasing and convex. If the agent remains unemployed at $t = 1$, he receives unemployment benefit b. If the agent successfully finds a job, he receives wage w at time $t = 1$ and pays a tax τ that is used to finance the UI system. Let $u(\cdot)$ denote the agent's utility over consumption at $t = 1$, where $u(\cdot)$ is strictly concave. Let $c_u = A + b$ and $c_e = A + w - \tau$ denote consumption if the agent remains unemployed or finds a job, respectively.

The social planner's problem is to choose the benefit level that maximizes the agent's expected utility subject to a budget balance constraint. In particular, the planner's problem is

$$\max_b \tilde{W}(b) = (1 - s(b))u(c_u) + s(b)u(c_e) - \psi(s(b))$$

such that

$$b(1 - s(b)) = s(b)\tau.$$

Taking the first order condition and applying the envelope theorem yields

$$\frac{d\tilde{W}}{db} = (1 - s)u'(c_u) - su'(c_e)\frac{d\tau}{db}.$$

Following Baily (1978) and Chetty (2006, 2008), consider the money metric marginal utility $\frac{dW}{db}$ obtained by scaling $\frac{d\tilde{W}}{db}$ by the welfare gain from a marginal increase in the wage, $su'(c_e)$. This yields

$$\frac{dW}{db} = \frac{d\tilde{W}/db}{su'(c_e)} = \frac{1-s}{s}\left\{\frac{u'(c_u) - u'(c_e)}{u'(c_e)} - \frac{\epsilon_{1-s,b}}{s}\right\} \tag{3.1}$$

where $\epsilon_{1-s,b}$ is the elasticity of the probability of remaining unemployed with respect to the benefit level, or the elasticity of one minus search effort s with respect to benefits b.

In general, this expression does not allow us to analyze the welfare impact of UI benefit changes unless we are willing to assume a particular form for $u(\cdot)$. Rather than choosing a functional form, following Baily (1978) and Chetty (2006) we approximate $u(\cdot)$ by its second-order Taylor expansion around c_e. In particular,

$$u'(c_u) - u'(c_e) \approx u''(c_e)(c_u - c_e)$$

which yields

$$\begin{aligned}\frac{u'(c_u) - u'(c_e)}{u'(c_e)} &\approx -\frac{u''(c_e)}{u'(c_e)}c_e\frac{c_e - c_u}{c_e} \\ &= \gamma\frac{c_e - c_u}{c_e} = \gamma\Delta \end{aligned} \tag{3.2}$$

where $\gamma = -\frac{u''(c_e)}{u'(c_e)}c_e$ is the coefficient of relative risk aversion evaluated at c_e and $\Delta = \frac{c_e - c_u}{c_e}$ is the

proportional consumption drop from unemployment.[5] We are left with the familiar Baily-Chetty formula:

$$\frac{dW}{db} \approx \frac{1-s}{s}\left\{\gamma\Delta - \frac{\epsilon_{1-s,b}}{s}\right\}. \tag{3.3}$$

This can be interpreted as the welfare change associated with raising the level of UI payments. Intuitively, the first bracketed term captures the value of transferring money from the agent's employed state to the agent's unemployed state, while the second term reflects the moral hazard cost of raising benefits: increasing benefits by a dollar increases the cost of the program by more than a dollar per unemployed agent because agents may respond to higher benefits by searching less intensely.

3.3 Welfare Analysis Under Heterogeneity

We are interested in extending the analysis above to contexts with heterogeneous agents. In particular, suppose that instead of a single unemployed agent there is a population of agents indexed by $i \in I$ for some index set I and that the distribution of agents is given by F. By the same logic as above, the marginal utility of agent i from an increase in the UI benefit level is $(1-s^i)u_i'(c_u^i)$, while the marginal disutility of a tax increase is $s^i u_i'(c_e^i)$. Thus, the welfare change for agent i from a \$1 increase in the UI benefit together with a tax change $\frac{d\tau}{db}^i$ in the employed state is

$$\frac{d\tilde{W}}{db}^i = \left(1-s^i\right)u_i'(c_u^i) - s^i u_i'(c_e^i)\frac{d\tau}{db}^i. \tag{3.1}$$

While the same approach as above can be used to approximate this marginal welfare for each agent, to measure the overall effect of a given change in the UI system we need some way to aggregate across individuals. Aggregation of individual welfare gains $\frac{d\tilde{W}}{db}^i$ is complicated by the well-known fact that Von-Neumann Morgenstern utilities are only defined up to an affine transformation: the behavior of an individual with utility function $u_i(\cdot)$ is in every way indistinguishable from that of an individual with utility function

$$\tilde{u}_i(\cdot) = a_1^i + a_2^i u_i(\cdot) \tag{3.2}$$

for a_1^i a real number and $a_2^i > 0$. Even if we take the utility function for individual i to have some cardinal meaning, that is meaning beyond merely representing preferences, we have no hope of recovering quantities that are sensitive to the constants a_1 and a_2 without additional assumptions. Hence, to obtain an empirically implementable expression for aggregate welfare we need either to ensure invariance to individual-specific affine transformations of utility (3.2) or to impose additional assumptions that rule out problematic transformations.

We see at least three ways to proceed. One is to normalize individual utilities to rule out problematic transformations of the form (3.2). Another is to choose welfare weights that, by construction,

[5] Chetty (2006) notes that if the third order term of $u(c)$ is significant, the approximation should include an additional term that depends on the coefficient of relative prudence, $\rho = -\frac{u'''(c)}{u''(c)}c$, an adjustment for precautionary saving motives. We do not include this term in our baseline analysis, but discuss it in section 3.6.

generate invariant measures of aggregate welfare. Finally rather than attempting to aggregate the marginal welfare gains (3.1) directly, we can instead consider quantities, for example money-metric utilities, that are more directly comparable across individuals and aggregate these. Below, we show that reasonable implementations of all three approaches lead to equivalent expressions for marginal aggregate welfare.

3.3.1 Normalization of Utilities

To rule out potentially problematic transformations of the form (3.2), we can normalize the utility function to eliminate a degree of freedom. In particular, note that (3.1) depends only on the marginal utility $u_i'(\cdot)$ and hence it suffices to rule out multiplication of the utility function by a constant a_2^i. One way to do this is to normalize marginal utility in the employed state to one, effectively fixing $a_2^i = u_i'(c_e^i)^{-1}$. Under this normalization, (3.1) becomes

$$\frac{d\tilde{W}^i}{db} = \left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}.$$

If we consider a utilitarian welfare metric, which takes aggregate welfare to be $\hat{W} = E\left[\tilde{W}^i\right]$, we have that the welfare change from the reform considered is

$$\frac{d\hat{W}}{db} = E\left[\frac{d\tilde{W}^i}{db}\right] = E\left[\left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}\right]$$

$$= (1 - \bar{s})\, E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + (1 - \bar{s}) - E\left[s^i \frac{d\tau^i}{db}\right] \tag{3.3}$$

where $E[x^i] = \bar{x} = \int_I x^i dF(i)$ denotes the mean of x^i in the population and $E^u[x^i] = \bar{x}^u = \int_I \frac{1-s^i}{1-\bar{s}} x^i dF(i)$ denotes the mean weighted by unemployment probability and is the average one obtains by considering a cross-section of unemployed agents at $t = 1$. Hence, under this normalization we see that marginal aggregate welfare depends on three terms: the unemployment rate at $t = 1$, $1 - \bar{s}$, the average proportional increase in marginal utility from unemployment among the unemployed population, $E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right]$, and the average expected tax increase, $E\left[s^i \frac{d\tau^i}{db}\right]$.

3.3.2 Stabilizing Welfare Weights

We could aggregate the marginal welfare gains (3.1) without imposing a normalization, but must then choose welfare weights carefully to arrive at an invariant measure of aggregate welfare. To this end, suppose we attach welfare weight α^i to individual i and are interested in measuring $\hat{W} = E\left[\alpha^i \tilde{W}^i\right]$. The change in this aggregate welfare measure from the UI reform considered is

$$\frac{d\hat{W}}{db} = E\left[\alpha^i \frac{d\tilde{W}^i}{db}\right] = E\left[\alpha^i \left(\left(1 - s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i) \frac{d\tau^i}{db}\right)\right]. \tag{3.4}$$

Note, however, that if $\alpha^i \left(\left(1 - s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i) \frac{d\tau}{db}^i \right)$ is non-invariant to the transformation (3.2) for a positive mass of agents then (3.4) is non-invariant as well. Hence, invariance of (3.4) implies that for almost every agent i we have $\alpha^i \frac{d\tilde{W}}{db}^i = \tilde{K}^i$ for a constant $\tilde{K}^i$ which is invariant to (3.2). Provided $\frac{d\tilde{W}}{db}^i \neq 0$ for almost every agent, we have that $\alpha^i = \tilde{K}^i / \frac{d\tilde{W}}{db}^i$ which can be seen to imply that $\alpha^i = \frac{K^i}{u_i'(c_e^i)}$, again for K^i invariant to (3.2). Hence, to construct an invariant marginal aggregate welfare (3.4) all that remains is to pick the constant K^i. The choice $K^i = 1$ yields

$$\alpha^i \frac{d\tilde{W}}{db}^i = \left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau}{db}^i$$

and hence that as in the previous section,

$$\frac{d\hat{W}}{db} = (1 - \bar{s}) E^u \left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)} \right] + (1 - \bar{s}) - E\left[s^i \frac{d\tau}{db}^i \right].$$

Note that this result depends on our choice of K^i; as might be expected, other choices of welfare weights yield different expressions for aggregate welfare.

3.3.3 Aggregating Money-Metric Utilities

A third approach to constructing invariant measures of aggregate benefit from a UI reform is to pick some non-utility measure of benefit that is more easily comparable across individuals and to aggregate that. In this section, we consider aggregating money-metric utilities.

The welfare gain for individual i from the considered UI and tax changes, relative to their welfare gain from a \$1 wage increase, is

$$\frac{dW^i}{db} = \frac{d\tilde{W}^i/db}{s^i u_i'(c_e^i)} = \frac{\left(1 - s^i\right) u_i'(c_u^i)}{s^i u_i'(c_e^i)} - \frac{d\tau}{db}^i.$$

Since agent i regains employment with probability i, the expected cost of delivering this increased wage to agent i in the employed state, holding search behavior fixed, is s^i. Hence, the expected cost, fixing search intensity s^i, of the combination of agent-specific wage increases that changes each agent's expected utility by the same amount as the proposed UI reform is

$$\begin{aligned} E\left[s^i \frac{dW^i}{db} \right] &= E\left[\frac{\left(1 - s^i\right) u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau}{db}^i \right] \\ &= (1 - \bar{s}) E^u \left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)} \right] + (1 - \bar{s}) - E\left[s^i \frac{d\tau}{db}^i \right] \end{aligned}$$

which is exactly the same as the welfare expressions derived above. Note that this also has the interpretation as the cost, holding search behavior fixed, of the combination of agent-specific wage changes such that all agents would be indifferent between this wage change and the proposed reform. The cost per employed agent, which is more closely analogous to the money-metric utility considered

in the single-agent case, is then

$$\frac{d\bar{W}}{db} = \frac{1-\bar{s}}{\bar{s}} E^u \left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)} \right] + \frac{1-\bar{s}}{\bar{s}} - E^e \left[\frac{d\tau^i}{db} \right] \tag{3.5}$$

where $E^e[x^i] = \bar{x}^e = \int_I \frac{s^i}{\bar{s}} x^i dF(i)$ denotes an average weighted by agents' job-finding probability. This expression depends on the same terms as (3.3) above but has a more intuitive interpretation. In particular it depends on the average increase in marginal utility from unemployment among the unemployed population, $E^u \left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)} \right]$, the number of unemployed agents per employed agent, $\frac{1-\bar{s}}{\bar{s}}$, and the average tax increase faced by employed workers, $E^e \left[\frac{d\tau^i}{db} \right]$. Hence, we can see that optimal UI policy based on this metric will balance the consumption smoothing benefit from UI for the unemployed against the marginal disutility of taxes for employed agents, taking into account the relative size of these two populations. This expression will be our focus for the remainder of the analysis, but by the results discussed above one could also multiply all our expressions by $\bar{s}$ and interpret our results as concerning a utilitarian welfare metric under a particular normalization or particular welfare weights.

3.3.4 Alternative Approaches to Evaluating Aggregate Welfare

Our analysis represents only one of many possible approaches to extending the Baily (1978) formula to models with heterogeneous agents. A natural alternative, implicit in Chetty (2006), is to instead consider the average marginal utility from a given UI and tax change relative to the average marginal utility of a \$1 wage increase, i.e.

$$\frac{dW^*}{db} = \frac{E\left[d\tilde{W}^i/db\right]}{E\left[s^i u_i'(c_e^i)\right]} = \frac{E\left[\left(1-s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i) \frac{d\tau^i}{db}\right]}{E\left[s^i u_i'(c_e^i)\right]}.$$

The marginal welfare $\frac{dW^*}{db}$ values a UI reform in terms of the wage increase that generates the same average welfare change. We can interpret this quantity as the uniform wage change such that a rational agent would be indifferent between this change and the proposed reform if they knew they would subsequently be assigned an identity i at random according to the population distribution. In contrast, $\frac{d\bar{W}}{db}$ calculates the wage increase equivalent to the reform for each agent i separately and then aggregates across the population. Analogous to the literature on efficiency in economics with incomplete information we can view $\frac{dW^*}{db}$ as an *ex-ante* money-metric measure of welfare from the proposed reform (see Holmstrom and Myerson, 1983), since in calculating both the utility gain from the proposed UI reform and the marginal utility from a wage increase we average across types i. The marginal welfare $\frac{d\bar{W}}{db}$ can instead be viewed as an *interim* money metric, since the money metric value of the reform for each agent i takes into account all relevant agent characteristics.

While both ex-ante and interim welfare criteria have their advantages, we argue that for the UI analysis considered here the interim money-metric $\frac{d\bar{W}}{db}$ has a number of important strengths. First, if one could survey a cross-section of agents and obtain truthful responses on the i-specific

value (in wage terms) of the proposed reform $\frac{d\bar{W}^i}{db}$ along with job-finding probability s^i this would suffice to calculate $\frac{d\bar{W}}{db}$ directly without any assumptions. Further, as we demonstrate in the next section even without such survey data obtaining approximations to $\frac{d\bar{W}}{db}$ for the heterogeneous case is straightforward, and the analysis readily accommodates extensions to more general contexts, for example benefits and taxes proportional to heterogeneous wages as discussed in section 3.6. While as noted in Chetty (2006) (footnote 8), if we take the tax change $\frac{d\tau}{db}^i$ to be uniform $\frac{d\tau}{db}^i = \frac{d\tau}{db}$ and assume that agents share a common coefficient of relative risk aversion $\gamma^i = \gamma$ we can use the analysis of that paper to approximate $\frac{dW^*}{db}$, given the empirical literature documenting preference heterogeneity in a wide range of settings the assumption of homogeneous risk aversion seems unappealing.[6] In contrast to the analysis for $\frac{d\bar{W}}{db}$, extending the approximations for $\frac{dW^*}{db}$ to accommodate unobserved preference heterogeneity is far from straightforward.

3.4 Baily-Chetty Under Heterogeneity

To extended the Baily-Chetty analysis to heterogeneous agents using marginal aggregate welfare (3.5), for a given collection of i-specific tax changes $\frac{d\tau}{db}^i$ we can apply the same second-order approximations to the utility function as in the single-agent case above. Since under our assumptions agents may exhibit different behavioral responses to changes in the level of the UI benefit, some changes to the UI system may generate net transfers across individuals in expectation and these transfers may matter for our welfare analysis. To separate these *transfer effects* from the consumption smoothing benefit of a given UI reform, we first calculate the marginal welfare gain from an increase in benefits in an actuarially-fair UI system that generates no transfers across agents in expectation and then turn to a UI reform with a uniform tax change $\frac{d\tau}{db}^i = \frac{d\tau}{db}$.

3.4.1 Actuarially Fair UI Under Heterogeneity

We begin by considering the case where taxes are actuarially fair within individual: each individual's taxes in the employed state reflect their personal search intensity in the unemployed state so there is expected budget balance for each individual, that is

$$b(1 - s^i(b)) = s^i(b)\tau^i.$$

This implies that there are no expected transfers between individuals. Under this restriction one can show that a \$1 increase in the UI benefit must be accompanied by a tax increase

$$\frac{d\tau^i}{db} = \frac{1 - s^i}{s^i}\left(1 + \frac{\epsilon^i_{1-s,b}}{s^i}\right)$$

[6]Note, moreover, that we still need to impose an appropriate normalization on the utility function to rule out transformations of the form (3.2).

on individual i in the employed state. To approximate the change in aggregate welfare (3.5) from this UI reform, we consider the same second-order approximation to the utility function as in the single-agent case, $u_i'(c_u) - u_i'(c_e) \approx u_i''(c_e)(c_u - c_e)$. The marginal aggregate welfare (3.5) from a change in the UI benefit level can then be approximated by

$$\begin{aligned}
\frac{d\bar{W}}{db} &\approx \frac{1-\bar{s}}{\bar{s}} E^u\left[\gamma^i \frac{c_e^i - c_u^i}{c_e^i}\right] + \frac{1-\bar{s}}{\bar{s}} - E^e\left[\frac{1-s^i}{s^i}\left(1 + \frac{\epsilon_{1-s,b}^i}{s^i}\right)\right] \\
&= \frac{1-\bar{s}}{\bar{s}} E^u\left[\gamma^i \Delta^i - \frac{\epsilon_{1-s,b}^i}{s^i}\right] \\
&= \frac{1-\bar{s}}{\bar{s}}\left(\bar{\gamma}^u \bar{\Delta}^u + \mathrm{cov}^u\left(\gamma^i, \Delta^i\right) - E^u\left[\frac{\epsilon_{1-s,b}^i}{s^i}\right]\right). \qquad (3.1)
\end{aligned}$$

where $\mathrm{cov}^u\left(x^i, y^i\right) = E^u[x^i y^i] - E^u[x^i]E^u[y^i]$ is defined analogously to $E^u[x^i]$ and corresponds to the covariance in a cross-section of the unemployed.

This expression for marginal welfare $\frac{d\bar{W}}{db}$ depends on three terms. The first term, $\bar{\gamma}^u \bar{\Delta}^u$, is the product of the average risk aversion in the cross-section of the unemployed and the average consumption drop. This term is analogous to the term $\gamma\Delta$ in the homogeneous case, but the weighting here is important. In particular, if more risk averse agents are more likely to be unemployed, this increases the value of raising the UI benefit all else equal. The second term, $\mathrm{cov}^u\left(\gamma^i, \Delta^i\right)$, captures the covariance between risk aversion and consumption drops in the cross-section of the unemployed, and reflects the fact that unemployment insurance is more valuable if more risk averse agents are subject to larger risks. We refer to this new term as the *covariance effect.* The last term in the expression, $E^u\left[\frac{\epsilon_{1-s,b}^i}{s^i}\right]$, measures the behavioral response to higher benefits, and depends on the joint distribution of individual-specific elasticities $\epsilon_{1-s,b}^i$ and job-finding rates s^i. The form of this term reflects our assumption of actuarial fairness, and as we discuss in the next section relaxing this restriction allows us to obtain a more tractable expression. Nonetheless, we can see that if the population of agents is completely homogeneous (3.1) simplifies to (3.3), confirming that our analysis generalizes the Baily-Chetty formula to heterogeneous agents.

3.4.2 UI with Uniform Taxes Under Heterogeneity

In the previous section we considered actuarially fair UI systems, which set the tax τ^i on each agent based on that agent's probability of remaining unemployed. We may also be interested in UI systems that are not actuarially fair within individual. In particular, the actuarially fair tax depends not only on the individual-level elasticity of unemployment probability with respect to UI benefits, $\epsilon_{1-s,b}^i$, but on the individual-level probability of regaining employment s^i. In contexts with unrestricted heterogeneity identifying the joint distribution of these objects, to say nothing of their values of each individual, poses a daunting challenge. In this section we consider a UI change that, while maintaining budget balance in the aggregate, need not be actuarially fair on an individual level and hence may generate net transfers across individuals through the UI system. In particular,

we consider UI systems such that the tax τ is uniform across individuals.

Transfers with Uniform Taxes

If we allow net transfers across individuals through the UI system, budget balance requires only that

$$b(1-\bar{s}(b)) = \bar{s}(b)\tau$$

which implies that

$$\frac{d\tau}{db} = \frac{1-\bar{s}}{\bar{s}}\left(1+\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right)$$

where $\epsilon_{1-\bar{s},b}$ is the elasticity of the unemployment rate in the second period (i.e. $1-\bar{s}$) with respect to the level of UI benefits. Hence, we can see that the money-metric value of the expected transfer to individual i due to a \$1 increase in the UI benefit together with a uniform change in τ is

$$\frac{d\tau^i}{db} - \frac{d\tau}{db} = \frac{1-s^i}{s^i}\left(1+\frac{\epsilon^i_{1-s,b}}{s^i}\right) - \frac{1-\bar{s}}{\bar{s}}\left(1+\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right).$$

Since the second term is constant across i, variation in this term across individuals is driven entirely by the behavior of the first term. In particular, individuals who are likely to remain unemployed (so $\frac{1-s^i}{s^i}$ is large) or who have a large behavioral response to increased benefits (so $\epsilon^i_{1-s,b}$ is large) will be subsidized by those whose probability of unemployment and behavioral response are smaller.

Welfare Analysis with Uniform Taxes

Again taking a second-order approximation to the utility function for each individual i, we can approximate the marginal welfare gain from an increase in the level of UI benefits together with a uniform change in the tax rate as

$$\begin{aligned}\frac{d\bar{W}}{db} &\approx \frac{1-\bar{s}}{\bar{s}}E^u\left[\gamma^i\frac{c^i_e - c^i_u}{c^i_e}\right] + \frac{1-\bar{s}}{\bar{s}} - E^e\left[\frac{1-\bar{s}}{\bar{s}}\left(1+\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right)\right] \\ &= \frac{1-\bar{s}}{\bar{s}}E^u\left[\gamma^i\Delta^i - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right] \\ &= \frac{1-\bar{s}}{\bar{s}}\left(\bar{\gamma}^u\bar{\Delta}^u + \mathrm{cov}^u\left(\gamma^i,\Delta^i\right) - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right). \end{aligned} \tag{3.2}$$

Hence, the welfare change from an increase in the UI benefit again depends on three terms: the product of the average risk aversion and average consumption drop in the unemployed population $\bar{\gamma}^u\bar{\Delta}^u$, the covariance between consumption drops and risk aversion $\mathrm{cov}^u\left(\gamma^i,\Delta^i\right)$ and, distinct from the actuarially fair case, $\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}$, the elasticity of the unemployment rate with respect to UI benefits divided by the average job finding probability.

If we consider a homogeneous population (3.2) again simplifies to the Baily-Chetty sufficient statistic (3.3) for the single agent case. If agents are heterogeneous but risk aversion and consumption drops are uncorrelated, the covariance term disappears and we have that for uniform

taxes

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}}\left(\bar{\gamma}^u \bar{\Delta}^u - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right),$$

which is a simple modification of the familiar Baily-Chetty formula where we have replaced each term by an analogous population quantity. To accommodate arbitrary heterogeneity we require only one additional term: the covariance between risk aversion and consumption drops in the unemployed population.

3.4.3 Implementing the Formulas

To apply these formulas we need estimates of four terms: the weighted average risk aversion, the weighted average consumption drop, the appropriate tax and transfer term, and the weighted covariance of risk aversion and consumption drops.[7] While the term $E^u\left[\frac{\epsilon^i_{1-s,b}}{s^i}\right]$ depends on individual-level elasticities and job-finding probabilities and may be quite difficult to calculate in practice, $\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}$ depends only on aggregate quantities and estimates are available in the literature (e.g. Meyer 1990).

The final term, the covariance of individual risk aversion and consumption drops, is novel and does not enter in the homogeneous case. To the best of our knowledge, this covariance has not been investigated empirically. Note that the sign of $\text{cov}^u\left(\gamma^i, \Delta^i\right)$ is *a priori* ambiguous. If agents were identical except for their risk preferences, we might expect that $\text{cov}^u\left(\gamma^i, \Delta^i\right) \leq 0$ given that agents with higher relative risk aversion value proportional consumption smoothing more. However, this basic intuition may break down given other plausible types of heterogeneity. For example, suppose poorer workers are both more risk averse (locally, e.g. risk aversion is declining in wealth) and less able to smooth consumption across employment states (e.g. have less savings, poorer access to credit markets, or are less able to substitute labor supply across household members). This can reverse the above intuition and generate $\text{cov}^u\left(\gamma^i, \Delta^i\right) \geq 0$. In general, the sign and magnitude of $\text{cov}^u\left(\gamma^i, \Delta^i\right)$ may depend on the context.

3.5 Calibrating the Covariance Effect

Our analysis above demonstrates that when agents are heterogeneous, the optimal benefit depends on the covariance between relative risk aversion and consumption drops across agents. This suggests that further study of the covariance effect is necessary before we can apply the Baily-Chetty approach with confidence. However, the approach already relies on several approximations–perhaps we can ignore the covariance term without introducing substantive bias to the welfare analysis. To get a sense of how large this covariance may be, we use data on observed consumption drops associated

[7] It is worth emphasizing that average relative risk aversion $\bar{\gamma}^u$ will not, in general, correspond to population estimates derived under a homogeneity assumption. For example, Cohen and Einav (2007) compute a 'back-of-the-envelope' absolute risk aversion estimate using population averages that is more than 6 times smaller than their estimate for mean absolute risk aversion derived from a structural model that allows for individual-level heterogeneity.

with unemployment and plausible population distributions of risk preferences to calibrate a simple model of private consumption smoothing decisions.

The model is two periods. In the first period, a set of heterogeneous workers draw implicit 'unemployment insurance' prices randomly from a common distribution. We can think of these prices as reflecting the effective price of consumption smoothing, whether by borrowing, saving, private insurance, informal insurance (e.g. risk sharing across households), spousal labor supply, or another mechanism. Given this price, each agent decides how much insurance to purchase. In the second period, each agent faces an unemployment shock with some probability, and consumes his available resources. In the absence of insurance, the shock reduces the resources available to consume.

We assume each agent has CRRA preferences over consumption, $u(c) = \frac{c^{1-\gamma_i}}{1-\gamma_i}$ with relative risk aversion γ_i. Except for their preferences over consumption, agents are identical *ex-ante*. We also assume that each agent would face the same consumption profile in the absence of insurance, c_e and c_u, and has the same probability of becoming unemployed, $1-s$. We calibrate the distribution of prices to rationalize the observed distribution of consumption drops given the distribution of risk preferences that we assume. The object of interest is the covariance between observed consumption drops for those who become unemployed and their relative risk aversion.

Mathematically, each agent solves the optimization problem

$$\max_{\alpha} su(c_e - \alpha p_i) + (1-s)u(c_u + (1-p_i)\alpha)$$

where $p_i \in (0,1)$ is the per-unit price of insurance faced by agent i and α is the amount of insurance purchased.

Taking the first order condition yields

$$\begin{aligned} psu'(c_e - \alpha p) &= (1-p)(1-s)u'(c_u + (1-p)\alpha) \\ \Rightarrow \Delta &= 1 - \left(\frac{ps}{(1-p)(1-s)}\right)^{-\frac{1}{\gamma}} \end{aligned}$$

where Δ is the consumption drop, or

$$\begin{aligned} \Delta_i &= \frac{c_e - \alpha p - (c_u + (1-p)\alpha)}{c_e - \alpha p} \\ &= \frac{c_e - c_u - \alpha}{c_e - \alpha p}. \end{aligned}$$

Intuitively, higher γ means an agent values consumption smoothing more and will make insurance purchases such that Δ is closer to zero (i.e. $|\Delta|$ is decreasing in γ). Hence, with implicit prices that tend to be actuarially unfair, this simple model should produce a negative covariance between γ and consumption drops across agents. Note that these predictions do not depend on the particular form of preferences used here. For example, applying a second-order approximation for a general

utility function $u(\cdot)$ as above yields

$$psu'(c_e - \alpha p) = (1-p)(1-s)u'(c_u + (1-p)\gamma)$$
$$\Rightarrow \Delta \approx \frac{1}{\gamma}\left(\frac{ps}{(1-p)(1-s)} - 1\right).$$

Again, it can be shown that $|\Delta|$ is decreasing in γ.

To measure the distribution of consumption drops, we use data from the Panel Study of Income Dynamics (PSID) from 1968-1993. The PSID surveys a sample of families, and includes information on household demographics, labor market outcomes, and consumption. To measure consumption, we follow Gruber (1997) and use self-reported household food consumption expenditures deflated using the CPI. This includes the amount usually spent on food both at home and away from home, and the value of food stamps used. The sample consists of all heads of household who are employed at time $t-1$ and unemployed at time t. Following the literature (e.g., Gruber 1997; Chetty and Sziedl 2007; Kroft and Notowidigdo 2011), we approximate proportional consumption drops by the change in log consumption from $t-1$ to t. Following Gruber (1997), we (a) drop observations where food consumption is imputed; and (b) drop observations with more than a three-fold change in total food consumption. Following Chetty and Sziedl (2007), we also exclude households that change in size between years. This yields a mean log consumption change of 8.2 log points. The estimated dispersion is substantial: the standard deviation of these measured log consumption changes is 42 log points.

Consumption data is typically measured with considerable error, and so will tend to overstate the variability of consumption drops. Ahmed et al. (2010) analyze a Canadian survey similar to the Consumer Expenditures Survey (CEX) and suggest that about 75% of variation in consumption data is due to measurement error. To correct for this measurement error, we take the true standard deviation of log consumption changes to be 25% of the measured variation, or about 10.5 log points.

Given that the typical UI replacement rate is around 50%, we normalize $c_e = 1$ and take $c_u = 0.50$. As in Chetty (2008), we take $s = 0.946$.

Given our uncertainty about the distribution of risk preferences in the population, we try a variety of distributions. For half of the calibrations, we assume that γ is uniformly distributed and for the other half, following Cohen and Einav (2007), we assume γ follows a lognormal distribution. We choose three ranges of values for each distribution type, corresponding to the 5th-95th percentile range when γ is lognormally distributed. The first range we try is 1 to 5, the range of values for γ typically seen in calibrations.[8] Motivated by the work of Chetty and Sziedl (2007), who argue that short-run consumption commitments can substantially amplify risk aversion over moderate stakes, we also try a larger range of 1 to 20.[9] In addition, we try an intermediate range of 1 to 10.

[8]For example, Gruber (1997) uses the range 1 to 4 and Kroft and Notowidigdo (2011) use the range 2 to 4. Chetty (2008) mentions that his results are consistent with $\gamma \approx 5$, while a more parametric approach in Chetty (2003) implies a γ of around 7. Both are estimated in the unemployment insurance context.

[9]While this dispersion is substantial, it is conservative compared to the distribution estimated in Cohen and Einav (2007). They find that, in the context of Israeli auto insurance, the standard deviation of absolute risk aversion is

Table 3.1: Covariance Term Calibration

Distribution:	Uniform			Lognormal		
Range:	1-5	1-10	1-20	1-5	1-10	1-20
Cov(γ, Δ)	-0.034	-0.107	-0.284	-0.036	-0.109	-0.272
$\bar{\gamma} \times \bar{\Delta}$	0.221	0.406	0.778	0.186	0.299	0.503
$\frac{\lvert Cov(\gamma,\Delta)\rvert}{\bar{\gamma}\times\bar{\Delta}}$	0.154	0.263	0.365	0.195	0.363	0.540
$\bar{\gamma}$	3.000	5.500	10.501	2.521	4.038	6.772
σ_γ	1.155	2.598	5.484	1.311	3.207	7.707
$\frac{\sigma_\gamma}{\bar{\gamma}}$	0.385	0.472	0.522	0.520	0.794	1.138
$\bar{p}$	0.067	0.075	0.090	0.064	0.066	0.068
σ_p	0.013	0.018	0.022	0.010	0.010	0.006

Notes: For lognormal distribution, the range is the 5th-95th percentile range. Results are based on 10 million i.i.d. draws for each case. The mean proportional consumption drop, $\bar{\Delta}$, is stable across cases at 7.4%.

We assume that prices are drawn i.i.d. from a normal distribution with mean μ_p and standard deviation σ_p, and then censor these prices at 0.01 and 0.99.[10] We find the values for μ_p and σ_p that match the mean and standard deviation of log consumption changes found in the PSID to those implied by the simulations.

Calibration results are presented in Table 3.1, and are based on 10 million i.i.d. draws for each case. There are two main results to note. First, and perhaps unsurprisingly, the larger the spread of γ relative to the mean, or $\frac{\sigma_\gamma}{\bar{\gamma}}$, the larger in magnitude the covariance term is relative to to the mean term, $\bar{\gamma} \times \bar{\Delta}$. Second, the covariance term is substantial in magnitude, ranging from 15% to 54% of the mean term, $\bar{\gamma} \times \bar{\Delta}$. Hence, our calibrations suggest that the covariance effect may have empirically significant implications for the Baily-Chetty approach to welfare analysis.

While the correlation between γ and Δ is negative in this simple example, recall that the sign of the covariance term in a more general model is ambiguous. In particular, the probability of an unemployment shock and the implicit price of consumption smoothing may vary arbitrarily with local risk preferences.

3.6 Extensions

In this section we consider three extensions to the analysis developed above. In the first subsection we extend our sufficient statistics approach to accommodate taxes and benefits that are proportional to wages, as in many real-world UI systems. Next, we follow Chetty (2006) and consider the impact

an order of magnitude larger than the mean. In the UI context, the empirical results in Chetty (2008) are consistent with substantial risk preference heterogeneity by (predicted) wealth quartile. Also see Barsky et al. (1997) who find evidence of substantial heterogeneity in implied risk aversion using survey responses.

[10]The estimates are generally insensitive to the censoring points chosen.

of third-order terms in the utility function, which Chetty argues may have a quantitatively important effect in some plausible cases. Finally, we show that using results from Chetty (2006) we can apply our approach in a rich class of dynamic models, allowing heterogeneity-robust sufficient statistics to be used in a far wider range of contexts than the simple static model discussed above.

3.6.1 UI with Proportional Benefits and Taxes

In our analysis above we consider a UI system with a constant benefit b and constant tax τ, consistent with the canonical Baily-Chetty model. In practice, however, UI benefits and taxes are often set proportional to individual wages, at least up to a cap. While this makes little difference in the single worker case, it affects the welfare analysis when workers are heterogeneous. In this section, we consider a UI system where benefits and taxes are proportional to the wage rate.

As a first step, again consider the single worker case. Let b_ρ and τ_ρ denote the UI replacement rate and tax rate. We now have that $c_u = A + b_\rho w$ and $c_e = A + w(1 - \tau_\rho)$. Consequently, the marginal welfare gain from an increase in b_ρ is

$$\frac{d\tilde{W}_\rho}{db_\rho} = w(1-s)u'(c_u) - wsu'(c_e)\frac{d\tau_\rho}{db_\rho}$$

and we arrive at the original Baily-Chetty formula after applying a modified scaling to obtain a money metric. In particular, instead of scaling by $su'(c_e)$, the welfare gain from a marginal increase in the wage, we scale by $wsu'(c_e)$, the welfare gain from a marginal *proportional* increase in the wage.[11]

As in section 3.3, to conduct welfare analysis for heterogeneous agents we need some way to meaningfully aggregate welfare. We will focus on the generalization of the money-metric approach discussed in section 3.3.3, but suitable extensions of the approaches discussed in sections 3.3.1 and 3.3.2 will yield equivalent analyses. To aggregate money metric utilities, note first that the welfare gain for individual i from a given proportional UI and tax change, relative to their welfare gain from a proportional increase in the wage, is

$$\frac{dW_\rho^i}{db_\rho} = \frac{d\tilde{W}_\rho^{\,i}/db_\rho}{s^i w^i u_i'(c_e^i)} = \frac{\left(1-s^i\right)u_i'(c_u^i)}{s^i u_i'(c_e^i)} - \frac{d\tau_\rho^{\,i}}{db_\rho}.$$

Since agent i regains employment with probability s^i and earns wage w^i while employed, the expected cost of delivering a proportional increase in agent i's wage, holding search behavior fixed, is $s^i w^i$. Hence, the expected cost of the combination of agent-specific proportional wage changes with

[11] Note that $\epsilon_{1-s,b_\rho} = \epsilon_{1-s,b}$.

the same effect on each agent's expected utility as the proposed UI change is

$$E\left[s^i w^i \frac{dW_\rho^i}{db_\rho}\right] = E\left[\frac{\left(w^i - s^i w^i\right) u_i'(c_u^i)}{u_i'(c_e^i)} - s^i w^i \frac{d\tau_\rho}{db_\rho}^i\right]$$
$$= \left(\bar{w} - \overline{sw}\right) E^{u,w}\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + \left(\bar{w} - \overline{sw}\right) - E\left[s^i w^i \frac{d\tau}{db}^i\right]$$

where $E^{u,w}[x^i] = \bar{x}^{u,w} = \int_I \frac{w^i - s^i w^i}{\bar{w} - \overline{sw}} x^i dF(i)$ denotes the average weighted by wage earnings lost due to unemployment. This is the average one obtains by considering a cross-section of unemployed agents at $t = 1$ weighted by individual wages. Hence the expected cost, relative to total wage income, for the combination of wage changes that would increase each individual's utility by the same amount as a one percentage point increase in the UI replacement rate together with individual-specific proportional tax changes $\frac{d\tau_\rho}{db_\rho}^i$ is

$$\frac{d\bar{W}_\rho}{db_\rho} = E\left[\frac{s^i w^i}{\overline{sw}} \frac{dW_\rho^i}{db_\rho}\right] = \frac{\bar{w} - \overline{sw}}{\overline{sw}} E^{u,w}\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right]$$
$$+ \frac{\bar{w} - \overline{sw}}{\overline{sw}} - E^{e,w}\left[\frac{d\tau_\rho}{db_\rho}^i\right] \tag{3.1}$$

where $E^{e,w}[x^i] = \int_I \frac{s^i w^i}{\overline{sw}} x^i dF(i)$ is defined analogously to $E^{u,w}[x^i]$. This term is closely related to (3.5), the primary difference being that all quantities in (3.1) depend on wages. In particular, we can see that the marginal increase in aggregate welfare from a given proportional UI reform depends on the average increase in marginal utility from unemployment weighted by lost wage earnings $E^{u,w}\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right]$, the ratio of wage earnings lost to unemployment to total wage earnings $\frac{\bar{w} - \overline{sw}}{\overline{sw}}$, and the increase in expected tax payments relative to total wage earnings $E^{e,w}\left[\frac{d\tau}{db}^i\right]$.

To obtain usable approximations to (3.1) we follow the same approach adopted in section 3.4. In particular, for actuarially fair proportional taxes the welfare change from a one percentage point increase in the replacement rate is

$$\frac{d\bar{W}_\rho}{db_\rho} \approx \frac{\bar{w} - \overline{sw}}{\overline{sw}}\left(\bar{\gamma}^{u,w}\bar{\Delta}^{u,w} + \mathrm{cov}^{u,w}\left(\gamma^i, \Delta^i\right) - E^{u,w}\left[\frac{\epsilon_{1-s,b_\rho}^i}{s^i}\right]\right) \tag{3.2}$$

where $\mathrm{cov}^{u,w}\left(x^i, y^i\right) = E^{u,w}[x^i y^i] - E^{u,w}[x^i]E^{u,w}[y^i]$ corresponds to the covariance in a cross-section of the unemployed weighted by individual wages. As we'd expect this expression is similar to (3.1), the corresponding formula for constant benefits and taxes. However, instead of weighting population averages by unemployment probability $1 - s^i$, we weight by expected wage earnings lost due to unemployment, $w^i - s^i w^i$.

As before, it is straightforward to extend our analysis to the uniform tax case. With uniform proportional taxes, budget balance requires that

$$b_\rho(\bar{w} - \overline{sw}(b_\rho)) = \overline{sw}(b_\rho)\tau_\rho$$

which implies that

$$\frac{d\tau_\rho}{db_\rho} = \frac{\bar{w} - \overline{sw}}{\overline{sw}} \left(1 + \epsilon_{\bar{w}-\overline{sw},b_\rho} \frac{\bar{w}}{\overline{sw}}\right)$$

where $\epsilon_{\bar{w}-\overline{sw},b}$ is the elasticity of lost wage earnings with respect to the replacement rate. Hence, for uniform proportional taxes the welfare gain associated with a one percentage point increase in the replacement rate is

$$\frac{d\bar{W}_\rho}{db_\rho} \approx \frac{\bar{w} - \overline{sw}}{\overline{sw}} \left(\bar{\gamma}^{u,w}\bar{\Delta}^{u,w} + \mathrm{cov}^{u,w}\left(\gamma^i, \Delta^i\right) - \epsilon_{\bar{w}-\overline{sw},b_\rho} \frac{\bar{w}}{\overline{sw}}\right). \tag{3.3}$$

This expression is similar to (3.2), the corresponding formula for a constant benefit and tax system. Again, to accommodate heterogeneity in wages we weight all averages by lost wage earnings rather than unemployment duration. In this case, the relevant elasticity is also that of lost wage earnings (with respect to the replacement rate) rather than unemployment probability, accounting for the impact of heterogeneous wages on the budget balance constraint.

In practice UI systems often have both constant and proportional components, for example setting benefits and taxes proportional to wages up to some cap, w^c, so $b^i = b_\rho \min\{w^i, w^c\}$ and $\tau^i = \tau_\rho \min\{w^i, w^c\}$. The results of this section can also be applied to this mixed case, provided we replace the wage w^i by the *taxable* wage $\tilde{w}^i = \min\{w^i, w^c\}$ in all expressions. The resulting welfare expression $\frac{d\bar{W}_\rho}{db_\rho}$ can again be viewed as a money-metric for welfare, though the interpretation is more involved.

3.6.2 Coefficient of Relative Prudence

In our baseline analysis we assume that the utility function of agent i is well-approximated by a two-term Taylor expansion, i.e. that

$$u_i'(c_u) - u_i'(c_e) \approx u_i''(c_e)(c_u - c_e).$$

As noted in Chetty (2006) this approximation may problematic if third-order terms in the utility function are large. In particular Chetty discusses a calibration based Gruber (1997) with CRRA utility and risk aversion ranging from 1 to 5, and shows that the formula (3.3) for the single agent case sometimes underestimates the optimal level of UI benefits by more than 30%. To correct this issue Chetty suggests including third-order terms in the approximation to $u(\cdot)$, considering

$$u'(c_u) - u'(c_e) \approx u''(c_e)(c_u - c_e) + \frac{1}{2}u'''(c_e)(c_u - c_e)^2 \tag{3.4}$$

for the single-agent case, and shows that this approximation yields a new approximation to dW/db that depends on the coefficient of relative prudence $\rho = \frac{u'''(c_e)}{u''(c_e)}c_e$, which accounts for precautionary savings motives. In Chetty's calibration the inclusion of this term reduces the maximal error in estimating the optimal benefit level to less than 4%, though as Chetty notes there are certainly examples where this distortion would be larger.

Our heterogeneity-robust approach can easily accommodate such third order terms. Taking the exact expression for marginal welfare (3.5) and applying the approximation (3.4) for each individual we obtain

$$\begin{aligned}\frac{d\bar{W}}{db} &\approx \frac{1-\bar{s}}{\bar{s}} E^u\left[\frac{u_i''\left(c_e^i\right)\left(c_u^i-c_e^i\right)+\frac{1}{2}u_i'''\left(c_e^i\right)\left(c_u^i-c_e^i\right)^2}{u_i'(c_e^i)}\right]+\frac{1-\bar{s}}{\bar{s}}-E^e\left[\frac{d\tau^i}{db}\right] \\ &= \frac{1-\bar{s}}{\bar{s}} E^u\left[\gamma^i\Delta^i+\frac{1}{2}\gamma^i\rho^i\left(\Delta^i\right)^2\right]+\frac{1-\bar{s}}{\bar{s}}-E^e\left[\frac{d\tau^i}{db}\right].\end{aligned}$$

Note that the addition of a third-order term in our approximation to u_i has no effect on the tax term $E^e\left[\frac{d\tau^i}{db}\right]$. Hence, for actuarially-fair individual taxes $\frac{d\tau^i}{db}=\frac{1-s^i}{s^i}\left(1+\frac{\epsilon_{1-s^i,b}}{s^i}\right)$ we have

$$\begin{aligned}\frac{d\bar{W}}{db} \approx & \frac{1-\bar{s}}{\bar{s}}\left\{E^u\left[\gamma^i\Delta^i+\frac{1}{2}\gamma^i\rho^i\left(\Delta^i\right)^2\right]-E^u\left[\frac{\epsilon_{1-s^i,b}}{s^i}\right]\right\} \\ = & \frac{1-\bar{s}}{\bar{s}}\left\{\bar{\gamma}^u\bar{\Delta}^u+\frac{1}{2}E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right]+\operatorname{cov}^u\left(\gamma^i,\Delta^i\right)\right. \\ & \left.+\frac{1}{2}\operatorname{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)-E^u\left[\frac{\epsilon_{1-s^i,b}}{s^i}\right]\right\}.\end{aligned}$$

This differs from the approximation (3.1) which neglects third-order terms in including the term $\frac{1}{2}\left(E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right]+\operatorname{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)\right)$. This term can be re-written as

$$\frac{1}{2}\left(\bar{\gamma}^u\bar{\rho}^u+\operatorname{cov}^u\left(\gamma^i,\rho^i\right)\right)\left(\left(\bar{\Delta}^u\right)^2+\operatorname{var}^u\left(\Delta^i\right)\right)+\frac{1}{2}\operatorname{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)$$

which makes it clear that the third-order approximation to $\frac{d\bar{W}}{db}$ depends on four terms that did not appear in the second-order approximation (3.1): the mean coefficient of relative prudence $\bar{\rho}^u$, the covariance between risk aversion and relative prudence $\operatorname{cov}^u\left(\gamma^i,\rho^i\right)$, the variance of consumption drops $\operatorname{var}^u\left(\Delta^i\right)$, and the covariance of the product $\gamma^i\rho^i$ with the squared consumption drop, $\operatorname{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)$, all in the population of unemployed agents. Many of these additional terms are fairly intuitive: $\bar{\rho}^u$ can be viewed as analogous to the parameter ρ in the single agent case while $\operatorname{cov}^u\left(\gamma^i,\rho^i\right)$ captures any co-movement between risk aversion and relative prudence. Our use of a third-order approximation to the utility implies a second-order approximation to the marginal utility $u_i'(\cdot)$, making the variance of consumption drops $\operatorname{var}^u\left(\Delta^i\right)$ potentially important for calculating expected marginal utility.

Analysis of the uniform-tax case with $\frac{d\tau^i}{db}=\frac{d\tau}{db}=\frac{1-\bar{s}}{\bar{s}}\left(1+\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right)$ proceeds along the same lines, yielding the same transfer terms as before and marginal welfare equal to

$$\begin{aligned}\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}}\Big\{&\bar{\gamma}^u\bar{\Delta}^u+\frac{1}{2}E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right]+\operatorname{cov}^u\left(\gamma^i,\Delta^i\right) \\ &+\frac{1}{2}\operatorname{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)-\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\Big\}\end{aligned}$$

where the term $\frac{1}{2}\left(E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right]+\text{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right)\right)$ can be re-written as above.

3.6.3 Sufficient Statistics for Dynamic Models

The central result of Chetty (2006) is that a simple sufficient-statistic formula analogous to those in the previous section continues to hold under reasonable conditions in a rich dynamic model with a single agent. As discussed in section 3.3.4 above, Chetty notes that this result can be extended to the case with heterogeneous agents provided risk aversion is homogeneous, though his approach corresponds to a different welfare metric than the one considered in this paper. In this section we argue that provided Chetty's assumptions hold for each agent i a natural generalization of our approach for static models allows us to extend our analysis to cover dynamic models with arbitrarily heterogeneous agents. Essentially, we show that under welfare metrics analogous to those considered in section 3.3 we can use Chetty's results to obtain approximate marginal welfare expressions for each agent and then aggregate across agents as before.

We consider a potentially heterogeneous version of the model studied by Chetty (2006). This model is considerably more elaborate than the simple static model described in section 3.2: we briefly introduce the key terms for our analysis but refer the interested reader to Chetty (2006) for a full exposition. As in Chetty's model, assume that time is continuous and that all agents i live from $t=0$ to $t=1$. The state of the world at time t is indexed by a state variable ω_t that follows some arbitrary stochastic process and whose unconditional distribution at time t is $F_t(\omega_t)$. Agents choose behavior at time t conditional on ω_t including consumption $c^i(t,\omega_t)$ and a vector of M other behaviors $x^i(t,\omega_t)$, for example, search effort and private insurance purchases. We assume that utility is time-separable and that the flow utility at time t is $u_i\left(c^i\left(t,\omega_t\right),x^i\left(t,\omega_t\right)\right)$ where we will typically suppress the dependence on $x^i(t,\omega_t)$. Agent i's employment status at time t is tracked by $\theta_t^i\left(t,\omega_t\right)$, which is equal to one if the agent is employed and zero if the agent is unemployed. Assume that the agents are subject to a budget constraint as in Chetty (2006), earn income w^i and pay tax τ^i in the employed state, and receive lump sum benefit b in the unemployed state. Chetty's model also allows the agent's activities x^i to generate income $f^i\left(x^i\left(t,\omega_t\right)\right)$, and accommodates the imposition of N other constraints on agent i's choices.

To show that the welfare analysis discussed above can be extended to this case, let us denote by $V^i\left(b,\tau\right)$ agent i's maximized expected utility, i.e.

$$V^i\left(b,\tau^i\right)=\max_{c^i,x^i}\int\int u_i\left(c^i\left(t,\omega_t\right),x^i\left(t,\omega_t\right)\right)dF_t(\omega_t)dt$$

subject to the constraints on the agent's choice set (see page 1887 in Chetty (2006) for a formal description of this optimization problem, where our case differs only allowing all terms to depend on the identity of agent i). Define D^i to be agent i's expected fraction of lifetime unemployed,

$$D^i=\int\int\left(1-\theta^i\left(t,\omega_t\right)\right)dF_t(\omega_t)dt.$$

Provided Assumptions 1-3 and 5 of Chetty (2006) hold for each agent i (which we'll assume to be the case), Lemma 1 in Chetty (2006) establishes that

$$\frac{dV^i}{db} = D \cdot Eu_i'(c_u) - \frac{d\tau^i}{db}(1-D) \cdot Eu_i'(c_e)$$

where

$$Eu_i'(c_e) = \frac{\int\int \theta^i(t,\omega_t)\, u_i'\left(c^i(t,\omega_t), x^i(t,\omega_t)\right) dF_t(\omega_t)dt}{\int\int \theta^i(t,\omega_t)\, dF_t(\omega_t)dt}$$

$$Eu_i'(c_u) = \frac{\int\int \left(1-\theta^i(t,\omega_t)\right) u_i'\left(c^i(t,\omega_t), x^i(t,\omega_t)\right) dF_t(\omega_t)dt}{\int\int \left(1-\theta^i(t,\omega_t)\right) dF_t(\omega_t)dt}$$

are agent i's marginal utility in the employed and unemployed states, respectively, averaging over both time and states of the world.

To extend our welfare analysis for heterogeneous agents to this case, note that we face the same difficulties aggregating welfare across heterogeneous agents as discussed in section 3.3. To overcome this challenge we adapt our approach for the static case, normalizing agent i's utility function by $\frac{1}{Eu_i'(c_e)}$ or, equivalently, attaching welfare weight $\alpha^i = \frac{1}{Eu_i'(c_e)}$ to agent i. The expected normalized marginal welfare of agent i is then

$$\frac{dW^i}{db} = D^i \cdot \frac{Eu_i'(c_u)}{Eu_i'(c_e)} - \frac{d\tau^i}{db}\left(1-D^i\right)$$

and marginal aggregate welfare (dividing through by the constant $1-\bar{D}$) is

$$\begin{aligned}\frac{d\bar{W}}{db} = \frac{1}{(1-\bar{D})}\int \frac{dW^i}{db}dF_i = &\frac{\bar{D}}{(1-\bar{D})}E^u\left[\frac{Eu_i'(c_u)-Eu_i'(c_e)}{Eu_i'(c_e)}\right] \\ &+ \frac{\bar{D}}{(1-\bar{D})} - E^e\left[\frac{d\tau^i}{db}\right]\end{aligned} \tag{3.5}$$

where $\bar{D} = \int D^i dF_i$ while

$$E^u\left[X^i\right] = \frac{1}{\bar{D}}\int D^i X^i dF_i$$
$$E^e\left[X^i\right] = \left(\frac{1}{1-\bar{D}}\right)\int \left(1-D^i\right) X^i dF_i$$

are averages over agents weighted by expected fraction of life spent unemployed and employed, respectively.

To give a money metric interpretation for (3.5) note that by equation (14) in Chetty (2006) the marginal utility of agent i with respect to a permanent increase in consumption in the employed state is $\left(1-D^i\right) Eu_i'(c_e)$. Note, further, that the cost of providing a permanent \$1 consumption increase to agent i in the employed state, holding the expected fraction of lifetime unemployed D^i fixed, is $1-D^i$. Hence, the aggregate welfare (3.5) has the interpretation as the total expected

cost (holding job-finding rates fixed) of the bundle of agent-specific consumption changes in the employed state that increase the expected utility of each agent by the same amount as the proposed UI reform, directly generalizing the money metric for the static model to this much richer dynamic case.[12]

Chetty's Lemma 2 shows that for each individual i, provided third and higher order terms in the utility function are small we can approximate average marginal utility in each state by the marginal utility of the average consumption in that state, i.e. $Eu_i'(c_e) \approx u_i'\left(\bar{c}_e^i\right)$ and $Eu_i'(c_u) \approx u_i'\left(\bar{c}_u^i\right)$ for

$$\bar{c}_e^i = \frac{\int\int \theta^i\left(t,\omega_t\right) c^i\left(t,\omega_t\right) dF_t(\omega_t)dt}{\int\int \theta^i\left(t,\omega_t\right) dF_t(\omega_t)dt}$$

$$\bar{c}_u^i = \frac{\int\int \left(1-\theta^i\left(t,\omega_t\right)\right) c^i\left(t,\omega_t\right) dF_t(\omega_t)dt}{\int\int \left(1-\theta^i\left(t,\omega_t\right)\right) dF_t(\omega_t)dt}$$

and a third-order expansion of the utility function $u_i(\cdot)$ then yields that

$$\frac{Eu_i'(c_u) - Eu_i'(c_e)}{Eu_i'(c_e)} \approx \gamma^i\Delta^i + \frac{1}{2}\gamma^i\rho^i\left(\Delta^i\right)^2$$

for $\Delta^i = \frac{\bar{c}_e^i - \bar{c}_u^i}{\bar{c}_e^i}$. Note that in computing Δ^i we use agent i's consumption in each employment state averaged over time and states of the world, so Δ^i can be interpreted as a particular measure of the drop in consumption for agent i due to unemployment. Substituting this approximation into (3.5) for actuarially fair tax change $\frac{d\tau}{db}^i = \frac{D^i}{1-D^i}\left(1+\frac{\epsilon_{D^i,b}}{1-D^i}\right)$ yields marginal aggregate welfare

$$\begin{aligned}\frac{d\bar{W}}{db} \approx \frac{\bar{D}}{(1-\bar{D})}\Bigg(&\bar{\gamma}^u\bar{\Delta}^u + \frac{1}{2}E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right] + \mathrm{cov}^u\left(\gamma^i,\Delta^i\right)\\ &+\frac{1}{2}\mathrm{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right) - E^u\left[\frac{\epsilon_{D^i,b}}{1-D^i}\right]\Bigg)\end{aligned} \tag{3.6}$$

for $\epsilon_{D^i,b}$ the elasticity of agent i's expected unemployment duration with respect to the benefit level. The expected money-metric value of the transfer to agent i under a uniform tax change $\frac{d\tau}{db} = \frac{\bar{D}}{1-\bar{D}}\left(1+\frac{\epsilon_{\bar{D},b}}{1-\bar{D}}\right)$ is $\frac{D^i}{1-D^i}\left(1+\frac{\epsilon_{D^i,b}}{1-D^i}\right) - \frac{\bar{D}}{1-\bar{D}}\left(1+\frac{\epsilon_{\bar{D},b}}{1-\bar{D}}\right)$, and the marginal aggregate welfare for a UI reform with uniform taxes $\frac{d\tau}{db}^i = \frac{d\tau}{db}$ is

$$\begin{aligned}\frac{d\bar{W}}{db} \approx \frac{\bar{D}}{(1-\bar{D})}\Bigg(&\bar{\gamma}^u\bar{\Delta}^u + \frac{1}{2}E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right]\\ &+\mathrm{cov}^u\left(\gamma^i,\Delta^i\right) + \frac{1}{2}\mathrm{cov}^u\left(\gamma^i\rho^i,\left(\Delta^i\right)^2\right) - \frac{\epsilon_{\bar{D},b}}{1-\bar{D}}\Bigg)\end{aligned} \tag{3.7}$$

where $\epsilon_{\bar{D},b}$ is the elasticity of the average unemployment duration with respect to the UI benefit level. Hence, we see that the third-order approximations obtained in section 3.6.2 above can be

[12]Note that in our baseline static model, a \$1 increase in the wage is equivalent to a \$1 consumption increase while employed.

generalized to a heterogeneous version of the dynamic model considered by Chetty (2006).

Note that we have followed the primary exposition in Chetty (2006) in using a second-order approximation to $u_i(\cdot)$ to approximate $Eu_i'(c_e) \approx u_i'\left(\bar{c}_e^i\right)$ but then using a third order expansion of $u_i(\cdot)$ evaluated at the mean consumption level. As in Chetty (2006) we could instead use a third order approximation to $u_i(\cdot)$ at both steps at the cost of introducing additional terms reflecting the relative variability of consumption in the employed and unemployed states. Alternatively, using a second-order approximation to $u_i(\cdot)$ at both steps yields (3.6) and (3.7) with $\rho^i \equiv 0$, and can be viewed as the generalization of our baseline approximation developed in section 3.4 to the dynamic case.

To understand (3.6) and (3.7) it is important to think carefully about precisely what averages are being taken. The expectation

$$E^u\left[X^i\right] = \frac{\int\int\int\left(1-\theta^i\left(t,\omega_t\right)\right)X^i dF_t(\omega_t)dtdF_i}{\int\int\int\left(1-\theta^i\left(t,\omega_t\right)\right)dF_t(\omega_t)dtdF_i}$$

takes the mean of X^i where agents are weighted by expected unemployment duration and corresponds to the mean of X^i that we would calculate by averaging over a random sample of agent-year observations in the unemployed state. Likewise, $\text{cov}^u\left(X^i, Y^i\right)$ is the covariance between X^i and Y^i that we would obtain by pooling repeated cross-sections of the unemployed population.

One aspect of the formulas (3.6) and (3.7) that could make them challenging to implement is the presence of the term Δ^i, which depends on the average consumption level for agent i across both time and ω_t conditional on each employment state. If we're willing to assume that agents face little lifetime consumption risk in each state, so that for each individual $\bar{c}_e^i$ and $\bar{c}_u^i$ are well-proxied by the average realized consumption in each state, then we can use any sample which contains the realized distribution of lifetime consumption paths and employment states to calculate $\bar{\Delta}^u$, $\text{var}^u\left(\Delta^i\right)$, and so on. Without such an assumption, however, the problem is more challenging and we will in general need additional assumptions (for example on the distribution of consumption risks faced by individuals) to recover the distribution of Δ^i from data on realized consumption profiles.

3.7 Conclusion

The Baily-Chetty formula is robust to a degree of heterogeneity but requires the assumption that agents share a common coefficient of relative risk aversion. In this paper, we extend the Baily-Chetty framework to allow for arbitrary heterogeneity across agents. We find that heterogeneity affects welfare analysis through the covariance effect: welfare gains depend on the covariance between risk aversion and consumption drops in the cross-section of the unemployed. This reflects the fact that unemployment insurance is more valuable if more risk averse agents are subject to larger risks. Calibration 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%.

Our results may have important implications for existing applications of the Baily-Chetty approach. For example, a recent literature extends the approach to investigate how UI benefits should vary over the business cycle (Landais et al. 2010; Kroft and Notowidigdo 2011; Schmieder et al. 2012). These papers emphasize that how optimal benefits vary over the business cycle depends on the cyclicality of the duration elasticity and consumption drops as a function of UI benefits. Our analysis makes clear that optimal benefits will also depend on the cyclicality of the covariance effect, which could arise if consumption smoothing mechanisms like consumer credit and spousal labor supply become less available during recessions.

Our results demonstrate that the value of social insurance depends on the covariance of risk exposure and risk aversion in the population: for a given distribution of consumption drops, the lower the covariance of risks faced by workers with individual risk aversion, the lower the value of additional social insurance. The sign and magnitude of the covariance effect are *a priori* ambiguous and will depend on the joint distribution of risk preferences, ex-ante risk, and ability to smooth consumption in a given context. Estimating this covariance is an important area for future research.

Bibliography

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4, Amsterdam: Elsevier-North Holland, 2011, pp. 1043–1171.

Ahmed, Naeem, Matthew Brzozowski, and Thomas F. Crossley, "Measurement Errors in Recall Food Consumption Data," April 2010. Unpublished manuscript.

Anderson, Bernard E., "The Ebb and Flow of Enforcing Executive Order 11246," *American Economic Review Papers and Proceedings*, May 1996, *86* (2), 298–301.

Arrow, Kenneth J., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, June 1962, *29* (3), 155–173.

Ashenfelter, Orley and James Heckman, "Measuring the Effect of an Antidiscrimination Program," in Orley Ashenfelter and James Blum, eds., *Estimating the labor market effects of social programs*, Princeton, NJ: Industrial Relations Section, 1976, chapter 3, pp. 46–84.

Aslund, Olof, Lena Hensvik, and Oskar Nordstrom Skans, "Seeking similarity: how immigrants and natives manage at the labor market," December 2009. IZA Discussion Paper Series No. 4640.

Athey, Susan, Christopher Avery, and Peter Zemsky, "Mentoring and Diversity," *American Economic Review*, September 2000, *90* (4), 765–786.

Autor, David, "Why Do Temporary Help Firms Provide Free General Skills Training?," *Quarterly Journal of Economics*, November 2003, *116* (4), 1409–1448.

Autor, David H., "The Economics of Labor Market Intermediation: An Analytic Framework," September 2008. NBER WOrking Paper No. 14348.

—— **and David Scarborough**, "Does Job Testing Harm Minority Workers? Evidence from Retail Establishments," *Quarterly Journal of Economics*, February 2008, *123* (1), 219–277.

Baily, Martin Neil, "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, December 1978, *10*, 379–402.

Barron, John M., Dan A. Black, and Mark A. Loewenstein, "Employer Size: The Implications for Search, Training, Capital Investment, Starting Wages, and Wage Growth," *Journal of Labor Economics*, January 1987, *5* (1), 76–89.

Barseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum, "Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data," *American Economic Review*, April 2011, *101* (2), 591–631.

Barsky, Roberet B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, May 1997, *112* (2), 537–579.

Barsky, Robert, John Bound, Kerwin Charles, and Joseph Lupton, "Accounting for the Black-White Wealth Gap: A Nonparametric Approach," *Journal of the American Statistical Association*, September 2002, *97* (459), 663–673.

Baum-Snow, Nathaniel, "Did Highways Cause Suburbanization?," *Quarterly Journal of Economics*, 2007, *122* (2), 775–805.

Beaman, Lori, Raghabendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova, "Powerful Women: Does Exposure Reduce Bias?," *Quarterly Journal of Economics*, November 2009, *124* (4), 1497–1540.

Blau, Francine D. and John W. Graham, "Black-White Differences in Wealth and Asset Composition," *Quarterly Journal of Economics*, 1990, *105* (2), 321–339.

Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy, and Jacque Eccles, "Empathy or Antipathy? The Impact of Diversity," *American Economic Review*, December 2006, *96* (5), 1890–1905.

Boustan, Leah, "Was Postwar Suburbanization "White Flight"? Evidence from the Black Migration," *Quarterly Journal*, February 2010, *125* (1), 417–443.

Boustan, Leah Platt and Robert A. Margo, "Race, Segregation, and Postal Employment: New Evidence on Spatial Mismatch," *Journal of Urban Economics*, 2009, *65*, 1–10.

Brown, Meta, Elizabeth Setren, and Giorgio Topa, "Do Informal Referrals Lead to Better Matches? Evidence from a Firm's Employee Referral System," August 2012. Federal Reserve Bank of New York Staff Report No. 568.

Card, David, "Immigrant Inflows, Native Outflows, and the Local Market Impacts of Higher Immigration," *Journal of Labor Economics*, January 2001, *19* (1), 22–64.

Carrington, William J., Kristin McCue, and Brooks Pierce, "Using Establishment Size to Measure the Impact of Title VII and Affirmative Action," *Journal of Human Resources*, Summer 2000, *35* (3), 503–523.

Chetty, Raj, "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," December 2003. Unpublished manuscript.

——, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90*, 1879–1901.

——, "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 2008, *116* (2), 173–234.

—— **and Adam Looney**, "Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies," *Journal of Public Economics*, December 2006, *90* (12), 2351–2356.

—— **and Adam Szeidl**, "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, May 2007, *122* (2), 831–877.

—— **and Emmanuel Saez**, "Optimal Taxation and Social Insurance with Endogenous Private Insurance," *American Economic Journal: Economic Policy*, 2010, *2* (2), 85–114.

Chung, Kim-Sau, "Role Models and Arguments for Affirmative Action," *American Economic Review*, June 2000, *90* (3), 640–648.

Coate, Stephen and Glenn C. Loury, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?," *American Economic Review*, December 1993, *83* (5), 1220–1240.

Cohen, Alma and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, June 2007, *97* (3), 745–788.

Cornell, Bradford and Ivo Welch, "Culture, Information, and Screening Discrimination," *Journal of Political Economy*, June 1996, *104* (3), 542–571.

Coulson, Edward, Derek Laing, and Paul Wang, "Spatial Mismatch in Search Equilibrium," *Journal of Labor Economics*, October 2001, *19* (4), 949–972.

Cutler, David M. and Edward L. Glaeser, "Are Ghettos Good or Bad?," *Quarterly Journal of Economics*, August 1997, *112* (3), 827–872.

Diamond, Peter A., "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," *American Economic Review*, 1967, *57* (4), 759–776.

Dobbin, Frank, *Inventing Equal Opportunity*, Princeton, NJ: Princeton University Press, 2009.

Donahue, John. J. and James Heckman, "Continuous Versus Episode Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature*, December 1991, *29* (4), 1603–1643.

Donahue, John J. and Peter Siegelman, "The Changing Nature of Employment Discrimination Litigation," *Stanford Law Review*, 1991, *43*, 983–1033.

Duraton, Giles and Matt Turner, "Urban Growth and Transformation," *Review of Economic Studies*. Forthcoming.

Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark Cullen. "How general are risk preferences? Choices under uncertainty in different domains," *American Economic Review*, October 2012, *102* (6), 2606–2638.

Fairlie, Robert W. and William A. Sundstrom, "The Emergence, Persistence, and Recent Widening of the Racial Unemployment Gap," *Industrial and Labor Relations Review*, January 1999, *52* (2), 252–270.

Fernandez, Roberto M. and Nancy Weinberg, "Sifting and Sorting: Personal Contacts and Hiring in a Retail Bank," *American Sociological Review*, December 1997, *62* (6), 883–902.

—— **, Emilio J. Castilla, and Paul Moore**, "Social Capital at Work: Networks and Employment at a Phone Center," *American Journal of Sociology*, March 2000, *105* (5), 1288–1356.

Finlay, Keith, "Effect of Employer Access to Criminal History Data on the Labor Market Outcomes of Ex-Offenders and Non-Offenders," in David H. Autor, ed., *Studies of Labor Market Intermediation*, Chicago: University of Chicago Press and NBER, 2009, pp. 89–125.

Fryer, Roland G. and Glenn C. Loury, "Valuing Diversity," *Journal of Political Economy*, 2013, *121* (4), 747–774.

___ **and Matthew O. Jackson**, "A Categorial Model of Cognition and Biased Decision Making," *The B.E. Journal of Theoretical Economics*, February 2008, *8* (1).

Gaines, Larry K., Patricia Costello, and Annis Crabtree, "Police Selection Testing: Balancing Legal Requirements and Employer Needs," *American Journal of Police*, 1989, *8*, 137–152.

General Accounting Office, *The Equal Employment Opportunity Program For Federal Nonconstruction Contractors Can Be Improved* April 1975.

___ , *Equal Employment Opportunity: DOL Contract Compliance Reviews Could Better Target Federal Contractors* September 1995.

Giuliano, Laura, David I. Levine, and Jonathan Leonard, "Manager Race and the Race of New Hires," *Journal of Labor Economics*, October 2009, *27* (4), 589–631.

Glaeser, Ed and Matthew Kahn, "Decentralized Employment and the Transformation of the American City," *Brookings-Wharton Papers on Urban Affairs*, 2001, *2*.

Goldstein, Morris and Robert S. Smith, "The Estimated Impact of Antidiscrimination Program Aimed at Federal Contractors," *Industrial Labor Relations Review*, July 1976, *29* (4), 523–543.

Gross, Tal and Matt Notowidigdo, "Health Insurance and the Consumer Bankruptcy Decision: Evidence from Expansions of Medicaid," *Journal of Public Economics*, 2011, *95* (7-8), 767–788.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, March 1997, *87* (1), 192–205.

Heckman, James J. and Kenneth I. Wolpin, "Does the Contract Compliance Program Work? An Analysis of Chicago Data," *Industrial and Labor Relations Review*, July 1976, *29* (4), 544–564.

Holmstrom, Bengt and Roger B. Myerson, "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, November 1983, *51* (6), 1799–1891.

Holzer, Harry and David Neumark, "Assessing Affirmative Action," *Journal of Economic Literature*, September 2000, *38* (3), 483–586.

___ **and** ___ , "What Does Affirmative Action Do?," *Industrial and Labor Relations Review*, January 2000, *53* (2), 240–271.

Holzer, Harry J., "Why Do Small Establishments Hire Fewer Blacks Than Larger Ones?," *Journal of Human Resources*, Autumn 1998, *33* (4), 896–914.

___ , **Steven Raphael, and Michael A. Stoll**, "Perceived Criminality, Criminal Background Checks, and the Racial Hiring Practices of Employers," *Journal of Law and Economics*, October 2006, *49* (2), 451–480.

Ihlanfeldt, Keith R. and David L. Sjoquist, "Job Accessibility and Racial Differences in Youth Unemployment Rates," *American Economics Review*, March 1990, *80*, 267–276.

Johnson, Rucker, "Landing a job in urban space: The Extent and effects of spatial mismatch," *Regional Science and Urban Economics*, 2006, *36*, 331–372.

Kain, John F., "Housing Segregation, Negro Employment and Metropolitan Decentralization," *Quarterly Journal of Economics*, 1968, *82* (3), 175–197.

Kalev, Alexandra, Frank Dobbin, and Erin Kelly, "Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies," *American Sociological Review*, August 2006, *71*, 589–617.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz, "Experimental Analysis of Neighborhood Effects," *Econometrica*, January 2007, *75* (1), 83–119.

Kroft, Kory and Matthew J. Notowidigdo, "Should Unemployment Insurance Vary With the Unemployment Rate? Theory and Evidence," December 2011. Working paper.

Kurtulus, Fidan Ana, "The Impact of Affirmative Action on the Employment of Minorities and Women Over Three Decades: 1973-2003," September 2011. Unpublished manuscript.

——, "Affirmative Action and the Occupational Advancement of Minorities and Women During 1973-2003," *Industrial Relations*, April 2012, *51* (2), 213–246.

Landais, Camille, Pascal Michaillat, and Emmanuel Saez, "Optimal Unemployment Insurance over the Business Cycle," November 2010. Working paper.

Lang, Kevin, "A Language Theory of Discrimination," *Quarterly Journal of Economics*, May 1986, *101* (2), 363–382.

Leonard, Jonathan S., "The Impact of Affirmative Action on Employment," *Journal of Labor Economics*, October 1984, *2* (4), 439–463.

——, "Affirmative Action as Earnings Redistribution: The Targeting of Compliance Reviews," *Journal of Labor Economics*, 1985, *3* (3), 363–384.

——, "What Promises Are Worth: The Impact of Affirmative Action Goals," *Journal of Human Resources*, Winter 1985, *20*, 3–20.

——, "Women and Affirmative Action," *Journal of Economic Perspectives*, Winter 1989, *3* (1), 61–75.

——, "The Impact of Affirmative Action Regulation and Equal Employment Law on Black Employment," *Journal of Economic Perspectives*, Autumn 1990, *4* (4), 47–63.

Lundberg, Shelly J. and Richard Startz, "Private Discrimination and Social Intervention in a Competitive Labor Market," *American Economic Review*, June 1983, *73* (3), 340–347.

Marsden, Peter V., "Selection Methods in US Establishments," *Acta Sociologica*, 1994, *37* (3), 287–301.

Masters, Adrian, "Matching with interviews," *Journal of Economic Dynamics and Control*, 2009, *33*, 938–954.

Mazzocco, Maurizio and Shiv Saini, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 2012, *102* (1), 428–468.

McCrary, Justin, "The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police," *American Economic Review*, 2007, *97* (1), 318–353.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook, "Birds of a Feather: Homophily in Social Networks," *Annual Review of Sociology*, August 2001, *27*, 415–444.

Meyer, Bruce, "Unemployment Insurance and Unemployment Spells," *Econometrica*, July 1990, *58*, 757–782.

Michaels, Guy, "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System," *Review of Economics and Statistics*, 2008, *90* (4), 683–701.

Miller, Amalia R. and Carmit Segal, "Does Temporary Affirmative Action Produce Persistent Effects? A Study of Black and Female Employment in Law Enforcement," *Review of Economics and Statistics*, 2012, *94* (4), 1107–1125.

Morgan, John and Felix Vardy, "Diversity in the Workplace," *American Economic Review*, March 2009, *99* (1), 472–482.

Mouw, Ted, "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*, 2000, *65*, 730–753.

Office of Federal Contract Compliance Programs, *Federal Contractor Compliance Manual* July 2013. Department of Labor.

Phelps, Edmund S., "The Statistical Theory of Racism and Sexism," *American Economic Review*, September 1972, *62* (4), 659–661.

Pincus, Fred L., *Reverse Discrimination: Dismantling the Myth*, Boulder, CO: Lynne Rienner Publishers, 2003.

Raphael, Steven, "The Spatial Mismatch Hypothesis and Black Youth Joblessness: Evidence from the San Francisco Bay Area," *Journal of Urban Economics*, 1998, *43*, 79–111.

Ritter, Joseph A. and Lowell J. Taylor, "Racial Disparity in Unemployment," *Review of Economics and Statistics*, February 2011, *93* (1), 30–42.

Rodgers, William M. and William E. Spriggs, "The Effect of Federal Contractor Status on Racial Differences in Establishment-Level Employment Shares: 1979-1992," *American Economic Review Papers and Proceedings*, May 1996, *86* (2), 290–293.

Schmieder, Johannes, Till von Wachter, and Stefan Bender, "The Effects of Extended Unemployment Insurance over the Business Cycle: Evidence from Regression Discontinuity Estimates over 20 Years," *Quarterly Journal of Economics*, 2012, *127* (2), 701–752.

Schulhofer-Wohl, Sam, "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, October 2011, *119* (5), 925–958.

Smith, James P. and Finis Welch, "Affirmative Action and Labor Markets," *Journal of Labor Economics*, 1984, *2* (2), 269–301.

Stiglitz, Joseph, "The Theory of 'Screening,' Education, and the Distribution of Income," *American Economic Review*, 1975, *65*, 283–300.

Weinberg, Bruce A., "Testing the Spatial Mismatch Hypothesis Using Inter-City Variations in Industrial Composition," *Regional Science and Urban Economics*, 2004, *34*, 505–532.

Wilson, Robert, "The Theory of Syndicates," *Econometrica*, 1968, *36* (1), 119–132.

Wilson, William Julius, *The Truly Disadvantaged: The Inner City, the Underclass and Public Policy*, Chicago, IL: University of Chicago Press, 1987.

——, *When Work Disappears: The World of the New Urban Poor*, New York, NY: Alfred A. Knopf, 1996.

Zax, Jeffery S. and John F. Kain, "Moving to the Suburbs: Do Relocating Companies Leave Their Black Employees Behind?," *Journal of Labor Economics*, 1996, *14* (3), 473–504.