Identifying and Estimating the Distributional Effects of Unionization and the Long-term Consequences of Military Service

by

Brigham R. Frandsen

B.S., Brigham Young University (2004)

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

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2010

Author ....
Department of Economics
August 13, 2010

Certified by ....
Joshua D. Angrist
Ford Professor of Economics
Thesis Supervisor

Certified by ....
Whitney K. Newey
Jane Berkowitz Carlton and Dennis William Carton Professor of Microeconomics
Thesis Supervisor

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

Accepted by ....
Esther Duflo
Abdul Latif Jameel Professor of Poverty Alleviation and Development Economics
Chairman, Departmental Committee on Graduate Studies

Identifying and Estimating the Distributional Effects of Unionization and the Long-term Consequences of Military Service

by

Brigham R. Frandsen

Submitted to the Department of Economics
on August 13, 2010, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis is concerned with the economic consequences for individuals of two important U.S. labor market institutions: unionization and the military draft.

The first chapter develops an econometric procedure for estimating quantile treatment effects in a regression discontinuity design. It shows nonparametric identification, develops estimators, including a data-driven bandwidth choice, and illustrates the methodology by estimating the effects of an Oklahoma universal pre-K program on the quantiles of student outcomes.

The second chapter applies the econometric procedure developed in the first chapter and estimates effects of unionization on the distribution of employees' earnings using a regression discontinuity design based on union certification elections. The results suggest that unionization raises the lower end of the distribution by up to 25 log points, but has a large negative effect on the upper tail of earnings, with little effect on average earnings. Unionization also increases retention among workers with lower pre-election earnings, but decreases it for higher-earning workers. These effects are interpreted as reflecting the political incentives unions face in certification elections.

The final chapter (joint with Joshua Angrist and Stacey Chen) explores the long-term effects of Vietnam-era military service on disability outcomes using a research design based on the draft lottery. We find no evidence that military service affected overall employment rates or overall work-limiting disability. At the same time, military service drastically increased federal transfer income, especially for lower skilled white men, among whom there was a large negative impact on employment and an increase in disability rates. The differential impact of Vietnam-era service on low-skilled men cannot be explained by more combat exposure for the least educated, leaving the relative attractiveness of VDC for less skilled men and the work disincentives embedded in the VDC system as a likely explanation.

Thesis Supervisor: Joshua D. Angrist
Title: Ford Professor of Economics

Thesis Supervisor: Whitney K. Newey
Title: Jane Berkowitz Carlton and Dennis William Carton Professor of Microeconomics

Thesis Supervisor: David Autor
Title: Professor of Economics

Acknowledgments

My name goes on the front of this thesis, but so many others have made it much better than it otherwise would have been.

First, Whitney Newey, David Autor, and Josh Angrist have been the paragon of advisers. I thank Whitney for his wise advice and unfailing encouragement. Thanks to David for asking (and demanding answers to) the tough questions. I thank Josh for never being satisfied and for hours of helpful conversation on the ski slopes or biking trails. My advisers have set the standard for what it means to be good economists, personally and professionally.

Besides my advisers many others have contributed more informally to making this thesis better. I thank Daron Acemoglu, Abhijit Banerjee, Victor Chernozhukov, Esther Duflo, Bob Gibbons, Jerry Hausman, Steve Pischke, and participants of the MIT Labor Lunch and Econometrics Lunch for their helpful suggestions through several iterations of this thesis.

Most of all, I thank my family. My wife, Christine, has shouldered the real burden of getting us through this, and my children, Elijah, Lucy, and Lincoln, have reminded me what really matters.

This thesis is dedicated to Christine.

Contents

Chapter 1

Nonparametric Identification and Estimation of Regression Discontinuity Quantile Treatment Effects

1.1 Introduction

The regression discontinuity (RD) design has received increased attention in recent years as a means of quasi-experimentally estimating treatment effects. To cite only a few examples of many recent studies using this design, Jacob and Lefgren (2004) and Matsudaira (2008) estimate the effect of remedial education programs, exploiting assessment test cutoffs in assignment to summer school programs; Black, Smith, Berger, and Noel (2003) use a feature of the Unemployment Insurance "profiling score" to evaluate the effect of the Worker Profiling and Reemployment Services program; Angrist and Lavy (1999) exploit maximum class size rules in Israeli public schools to estimate the effect of class size on educational outcomes; and DiNardo and Lee (2004) use certification elections to estimate the impact of new unions on employers.

The recent popularity of the RD design appears to be justified in many cases. Black, Galdo, and Smith (2007) and Buddelmeyer and Skoufias (2003) compare RD to randomized experiments and find the RD estimates replicate the experimental results well (see Cook and Wong, 2008, for a summary of studies comparing RD to experiments).

The studies mentioned above and others using the RD design focus on estimating local average treatment effects (Imbens and Angrist, 1994). In many contexts, however, the effect of a treatment on the entire distribution of outcomes is of interest. For example, economists often evaluate the social welfare implications of a policy based on the differences in the distribution of outcomes under various alternatives (Atkinson, 1970). Furthermore, a zero average effect may mask significant offsetting effects at different points in the distribution. Examples where distributional effects may be of particular interest include unionization, which is widely believed to compress wages, and progressively oriented social and education programs which may be intended to bring up the lower end of the distribution.

In this paper I introduce a procedure to estimate quantile treatment effects in the fuzzy RD design when selection into treatment is potentially endogenous. The results in this paper apply equally well to the sharp regression discontinuity design, which I treat as a special case of the fuzzy design. As Hahn, Todd, and van der Klaauw (2001) suggested, the fuzzy RD design leads naturally to instrumental variables (IV) type estimators, and the estimator they develop has an interpretation as a local Wald estimator of a local average treatment effect (LATE). Their insight suggests applying IV quantile treatment effects estimators in order to estimate distributional effects in the RD design.

Two recently developed approaches to IV quantile treatment effects are Chernozhukov and Hansen (2005), and Abadie, Angrist, and Imbens (2002) (see Frandsen (2008a) for a comparison of these two estimators). These two approaches rely on distinct sets of identifying assumptions, and the interpretations of the estimands differ. An RD quantile treatment effects estimator in the spirit of Chernozhukov and Hansen (2005) is developed by Guiteras (2008). In some contexts, however, the requirement of rank invariance or rank similarity across treatment status in that model may be less desirable than the LATE assumptions of Abadie, Angrist, and Imbens (2002). I therefore focus on the LATE framework in this paper. A challenge that prevents the trivial application of Abadie, Angrist and Imbens' (AAI) local quantile treatment effects estimator to the RD design is the fact that the instrument—an indicator for exceeding some threshold of a running (or forcing) variable—is a deterministic function of an included covariate, since the running variable must be controlled for. Thus conditional on included variables there is no variation in the instrument, and AAI's estimator is not defined.[1] One way to deal with this is to use kernel weighting to estimate the effect only at the threshold. In the limit the running variable plays no role, and therefore can be ignored, as in Froelich and Melly (2008). In finite sample, however, ignoring the running variable leads to substantial bias in this approach. The contribution of this paper is to overcome this difficulty using local linear quantile regression to nonparametrically estimate quantile treatment effects at the threshold, in the spirit of Hahn, Todd, and van der Klaauw (2001) and Porter (2003).

Another approach to estimating distributional effects in the RD context that makes use of local linear quantile regression is being developed by Baker, Firpo, and Milligan (2005). While their approach overcomes the finite sample sample bias problem inherent in the "local constant" approach of simply applying AAI at the threshold, they rely on a selection-on-observables identifying assumption at the threshold. This rules out cases where selection into treatment is endogenous, even at the threshold of the running variable. The estimator I introduce here allows for endogenous treatment even conditional on being in a neighborhood of the threshold, and thus it has an IV interpretation.

The remainder of this paper is outlined as follows. Section 1.2 develops the statistical framework. I establish identification results in Section 1.3, and I describe the estimation procedure in Section 1.4. I derive the asymptotic distribution for the estimator in Section 1.5, and discuss inference in Section 1.6. I present Monte Carlo simulation results in Section 1.7. In Section 1.8 I apply the procedure to estimate the effect of an Oklahoma universal pre-K program on the distribution of test scores, and Section 1.9 concludes.

1.2 Econometric framework

Since the motivation for the estimation procedure I develop in this paper is very much in the spirit of Imbens and Angrist's (1994) LATE framework, I will set up the fuzzy RD framework in terms of potential outcomes. For simplicity, initially I do not condition on any covariates other than the running variable, although I discuss including additional covariates in Section 1.4.2. The critical elements of the fuzzy RD design, in terms of the LATE notation, are:

$Y_0 \equiv$ potential outcome when untreated

[1] See Section 1.4 for more details on why AAI's approach fails in the RD context.

$Y_1 \equiv$ potential outcome when treated

$D \equiv$ indicator for treatment status (possibly endogenous)

$Y \equiv Y_0 + (Y_1 - Y_0) D$, observed outcome

$\delta \equiv Y_1 - Y_0$, treatment effect (possibly heterogeneous)

$R \equiv$ scalar running variable

$Z \equiv 1(R > 0)$, indicator for the running variable exceeding a threshold. I set the threshold equal to zero without loss of generality

$D_0 \equiv$ potential treatment status when $Z = 0$

$D_1 \equiv$ potential treatment status when $Z = 1$.

Some features this setup preserves from the LATE framework are that it allows for heterogeneous treatment effects and endogenous treatment selection, as in a Roy model of selection on gains. Another feature this setup shares with the LATE framework is that in a neighborhood around the threshold of the running variable, we can conceptually classify individuals into one of several mutually exclusive groups, depending on their potential treatment status. I will use the standard nomenclature for these groups, and introduce abbreviations to refer to them. Formally, I define these groups as events in a common probability space $(\Omega, \mathcal{F}, P)$:

- Always takers: $AT = \{\omega : D_1(\omega) = D_0(\omega) = 1\}$
- Never takers: $NT = \{\omega : D_1(\omega) = D_0(\omega) = 0\}$
- Compliers: $C = \{\omega : D_1(\omega) > D_0(\omega)\}$
- Defiers: $DE = \{\omega : D_1(\omega) < D_0(\omega)\}$.

The estimand I consider in this paper is the local quantile treatment effect, or the difference between the marginal distributions of potential outcomes for compliers at a particular quantile at the threshold level of the running variable:

$$\delta_{LQTE}(\tau) \equiv Q_{Y_1|C,R=0}(\tau) - Q_{Y_0|C,R=0}(\tau). \tag{1.1}$$

An important comment regarding the interpretation of this object is that it reflects the effect of treatment on the distribution, rather than the effect of treatment on any particular individual. Without a rank invariance assumption, as in the Chernozhukov and Hansen (2005) framework, there is no sense in which (2.2) represents the treatment effect for a particular individual, since an individual with a Y_0 of rank τ need not have a Y_1 of rank τ.

1.3 Identification of LQTE

Besides those embodied in the notation given in Section 1.2 , I make the following additional assumptions:

Assumption 1: RD $\lim_{r \longrightarrow 0^+} \Pr(D = 1|R = r) > \lim_{r \longrightarrow 0^-} \Pr(D = 1|R = r)$

Assumption 2: Local Smoothness $F_{Y_d|D_0,D_1,R}(y|d_0, d_1, r)$ is continuous in r over an ε-neighborhood of zero, and is strictly increasing in y over the same neighborhood, for $d \in \{0,1\}$. $E[D_d|R = r]$ is continuous in r for $r < 0$ or $r > 0$ in the same neighborhood, for $d \in \{0,1\}$.

Assumption 3: Monotonicity $\lim_{r \longrightarrow 0} \Pr(D_1 \geq D_0|R = r) = 1$

The first assumption is the defining feature of the regression discontinuity design, that the probability of treatment changes discontinuously at the threshold value of the running variable. Without loss of generality I assume the probability of treatment is greater above the threshold. Assumption 2 is a smoothness condition which, intuitively speaking, ensures that after controlling smoothly for the running variable, differences in the distribution of outcomes on either side of the threshold are due to the change in probability of treatment assumed in Assumption 1. Assumption 2 also guarantees quantiles of the potential outcomes are uniquely defined at the threshold. Assumption 3 is the crucial monotonicity assumption that the response of treatment selection to the instrument is monotone.[2] An immediate consequence of this assumption is that the monotonicity condition rules out the existence of defiers—those for whom $D_0 > D_1$—in a neighborhood around the threshold.

These assumptions are quite similar to Hahn, Todd, and van der Klaauw's (2001) conditions for identifying the local average treatment effect in an RD setting. Assumption

[2] There are several settings in which monotonicity holds automatically, including when non-compliance is one-sided, with either no treatment below the threshold, or 100 percent treatment above the threshold. Other settings which imply monotonicity are latent index models of selection, as discussed below.

1 here is precisely their RD condition, and Assumption 3 is equivalent to the monotonicity condition in their assumption A3. The smoothness of $F_{Y_d|R}(y|r)$, $\Pr(NT|R=r)$ and $\Pr(AT|R=r)$ in Assumption 2 are analogous to their assumption A1 and the joint independence condition in their A3. One difference is that Hahn, Todd, and van der Klaauw assume only the smoothness of the conditional expectation of $Y_1 - Y_0$, while I require smoothness of the conditional distribution function of Y_0 and Y_1 because I am identifying distributional effects, and the difference in quantiles is not the quantile of the difference.

The assumptions I make are analogous to those required for Abadie, Angrist, and Imbens's (2002) local quantile treatment effects estimator, or Imbens and Angrist's (1994) LATE assumptions. Instead of independence between an instrument and potential outcomes and potential treatment status, I make continuity assumptions on the distribution of potential outcomes and potential treatment status. The LATE first stage assumption is replaced by the analogous RD assumption that the probability of treatment jumps discretely as the running variable hits the threshold value. Assumption 3, local monotonicity, is directly analogous to the monotonicity assumption in the LATE framework. The most striking difference between my assumptions and Abadie, Angrist, and Imbens's (2002) assumptions is the absence here of the "Non-trivial assignment" condition which they require. Indeed, the principal challenge of applying Abadie, Angrist, and Imbens's (2002) quantile treatment effects estimator in an RD setting is that the non-trivial assignment condition fails here, since conditional on the running variable, the "instrument", $Z = 1(R > 0)$, is deterministically either zero or one.

Given our assumptions, at the threshold we can adapt Imbens and Rubin's (1997) and Abadie's (2002) method of identifying counterfactual distributions for compliers. The local quantile treatment effect is then simply the difference between the inferred marginal distributions of the potential outcomes for compliers at a particular quantile.

The following theorem, which is similar to that obtained by Froelich (2007) for the case with covariates, shows that the local quantile treatment effect can be written as the horizontal difference between "local Wald ratios", emphasizing the connection with instrumental variables estimation of treatment effects:

Theorem 1 *LQTE Identification.*

Under Assumptions 1-3, the local quantile treatment effect, (2.2), is identified from the *joint distribution of* (Y, D, R) *as*

$$\delta_{LQTE}(\tau) = F^{-1}_{Y_1|C,R=0}(\tau) - F^{-1}_{Y_0|C,R=0}(\tau),$$

where $F_{Y_1|C,R=0}(y)$ *and* $F_{Y_0|C,R=0}(y)$ *are given by*

$$F_{Y_1|C,R=0}(y) = \frac{\lim_{r\to 0^+} E[1(Y \le y) D|R=r] - \lim_{r\to 0^-} E[1(Y \le y) D|R=r]}{\lim_{r\to 0^+} E[D|R=r] - \lim_{r\to 0^-} E[D|R=r]} \quad (1.2)$$

$$F_{Y_0|C,R=0}(y) = \frac{\lim_{r\to 0^+} E[1(Y \le y)(1-D)|R=r] - \lim_{r\to 0^-} E[1(Y \le y)(1-D)|R=r]}{\lim_{r\to 0^+} E[1-D|R=r] - \lim_{r\to 0^-} E[1-D|R=r]} \quad (1.3)$$

Proof. All proofs are given in the Appendix. ■

As the theorem makes clear, not only are local quantile treatment effects identified, but the entire distributions of both potential outcomes conditional on compliers at the threshold are identified. Thus any parameter that is a function of the distributions of potential outcomes is also identified, including the counterfactual densities, distribution treatment effects, and measures of stochastic dominance.

1.4 Estimation procedure

The local quantile treatment effect, (2.2), may be consistently estimated in a number of ways. I will briefly discuss two approaches that may appear to be the most obvious, but I will spend greater time developing the preferred approach of local linear distribution function estimation.

As I suggested in the introduction, at first blush perhaps the most obvious approach to estimating the local quantile treatment effect is an adaptation of Abadie, Angrist, and Imbens's (2002) IV quantile treatment effects estimator. This estimator combines kernel weights which narrow in on the threshold with "complier finding" weights:

$$\kappa_v = 1 - \frac{D(1 - E[Z|Y,D])}{1 - E[Z]} - \frac{(1-D)E[Z|Y,D]}{E[Z]},$$

where $Z = 1(R \ge 0)$. These weights can be estimated to form $\hat{\kappa}_{vi}$, and the estimator would take the form:

$$\left(\hat{a}, \hat{\delta}_{\mathrm{ALQTE}}(\tau)\right) = \arg\min_{a,d} \frac{1}{n} \sum_{i=1}^{n} \rho_\tau \left(Y_i - a - dD_i\right) \cdot \hat{\kappa}_{vi} \cdot K\left(\frac{R_i}{h_n}\right), \tag{1.4}$$

where $K(\cdot)$ is a kernel function, $\rho_\tau(u) = (\tau - 1(u < 0))\,u$ and h_n is a smoothing parameter. Strictly speaking, all expectations in the definition of κ_v should be taken conditional on the running variable, R. In that case a zero would appear in one of the denominators, since Z is a deterministic function of R, violating the so-called non-trivial assignment condition which requires that there be variation in the instrument conditional on included variables. However, as the sample size grows, in the limit we are conditioning on $R = 0$, and so there is no longer any need to include R as a regressor, and thus the technique can be applied. However, in finite sample, there will be variation in R over the window defined by the kernel weights. This leads to a sort of omitted variables bias, since the instrument, Z, may not be independent of the potential outcomes when we do not condition on R. This finite sample bias is analogous to the inflated bias of locally constant kernel regression estimators, noted by Fan (1992). For the case of quantile regression, Yu and Jones (1997) show the bias in this approach is proportional to the slope of the conditional quantile function estimated. Another source of bias in this approach stems from implicitly estimating quantile functions at a boundary,[3] leading to a convergence rate of $n^{1/5}$ instead of the optimal nonparametric rate, $n^{2/5}$. The Monte Carlo results below show the combined bias from ignoring the running variable and boundary effects can be quite large.

An approach which overcomes these finite sample problems involves estimating quantiles of the potential outcomes for compliers at the discontinuity threshold. The key task is non-parametric quantile estimation at a boundary, an estimation problem considered by Yu and Jones (1998). Their preferred technique, which I adapt here, estimates conditional quantiles by inverting local linear estimates of the conditional distribution function.[4] The local linear technique is not subject to bias from ignoring the running variable and automatically corrects for boundary effects. (Fan, 1992)

[3]Thanks go to an anonymous referee for pointing this out.

[4]Yu and Jones (1998) also consider an estimator which minimizes a local linear "check function" to find the conditional quantiles directly, rather than indirectly via inverted distribution functions. I develop an estimator along these lines elsewhere (Frandsen, 2008). In practice it performs nearly identically to the inverted distribution function approach, but it takes longer to compute, so I focus on the preferred approach of inverting local linear estimates of distribution functions.

Given differentiability of the conditional distributions in a neighborhood of the threshold, a consistent estimator for the local quantile treatment effect, then, is the (horizontal) difference between local linear estimates of the conditional distribution functions (1.2) and (1.3) at a particular quantile:

$$\hat{\delta}_{LQTE}(\tau) = \hat{F}^{-1}_{Y_1|C,R=0}(\tau) - \hat{F}^{-1}_{Y_0|C,R=0}(\tau), \tag{1.5}$$

where

$$\begin{aligned} \hat{F}^{-1}_{Y_1|C,R=0}(\tau) &= \inf\left\{a : \hat{F}_{Y_1|C,R=0}(a) \geq \tau\right\}, & (1.6) \\ \hat{F}^{-1}_{Y_0|C,R=0}(\tau) &= \inf\left\{b : \hat{F}_{Y_0|C,R=0}(b) \geq \tau\right\}, & (1.7) \end{aligned}$$

and $\hat{F}_{Y_1|C,R=0}(y)$, $\hat{F}_{Y_0|C,R=0}(y)$ are local linear, consistent estimates of (1.2) and (1.3).

Since using local linear techniques to estimate (1.2) and (1.3) is the main insight, I will discuss possible ways that might be done. One approach is to recognize that (1.2) and (1.3) are "local Wald ratios" and estimate these quantities in one step via local linear two-stage least squares, as suggested by Imbens and Lemieux (2008) for mean estimation. However, an alternative approach is to recognize that (1.2) and (1.3) can be rewritten as:

$$F_{Y_1|C,R=0}(y) = \frac{\lim\limits_{r\to0^+} F_{Y|D=1,R=r}(y) \lim\limits_{r\to0^+} E[D|R=r] - \lim\limits_{r\to0^-} F_{Y|D=1,R=r}(y) \lim\limits_{r\to0^-} E[D|R=r]}{\lim\limits_{r\to0^+} E[D|R=r] - \lim\limits_{r\to0^-} E[D|R=r]} \tag{1.8}$$

$$F_{Y_0|C,R=0}(y) = \frac{\lim\limits_{r\to0^+} F_{Y|D=0,R=r}(y) \lim\limits_{r\to0^+} E[1-D|R=r] - \lim\limits_{r\to0^-} F_{Y|D=0,R=r}(y) \lim\limits_{r\to0^-} E[1-D|R=r]}{\lim\limits_{r\to0^+} E[1-D|R=r] - \lim\limits_{r\to0^-} E[1-D|R=r]} \tag{1.9}$$

The local linear two-stage least squares approach implicitly uses the same bandwidth to estimate the component quantities in (1.8) and (1.9), but in practice it is often optimal to use different bandwidths to estimate the various pieces, especially when the design is not balanced across the threshold.[5] The approach I propose estimates the pieces separately, allowing optimal choice of bandwidth for each piece. The estimator is therefore a function of first step estimates of the component quantities in (1.8) and (1.9), consisting of the

[5] An example of an unbalanced design is one-sided non-compliance where, for example, there are few or no treated units below the threshold. Many RD studies exhibit this, including Gormley, Gayer, Phillips, and Dawson (2005), which I use as an application in this paper, and DiNardo and Lee (2004).

following four conditional distributions:

$$\lim_{r\to 0^-} F_{Y|D=0,R=r}(y),$$
$$\lim_{r\to 0^+} F_{Y|D=0,R=r}(y),$$
$$\lim_{r\to 0^-} F_{Y|D=1,R=r}(y),$$
$$\lim_{r\to 0^+} F_{Y|D=1,R=r}(y)$$

and the following two conditional expectations:

$$\lim_{r\to 0^+} E[D|R=r] \tag{1.10a}$$
$$\lim_{r\to 0^-} E[D|R=r]. \tag{1.10b}$$

Since each of these quantities must be estimated at a boundary, local linear approaches are most suitable (Fan, 1992).[6] For the conditional distributions, possibly the most straightforward local linear estimator of $\lim_{r\to 0^-} F_{Y|D=0,R=r}(y)$, for example, would be $\hat{F}_{Y|D=0,R=0^-}(y)$ that satisfies:

$$\left(\hat{F}_{Y|D=0,R=0^-}(y), \hat{b}(y)\right) = \arg\min_{a,b} \sum_{i:D=0,Z=0} \left[1(Y_i \le y) - a - bR_i\right]^2 K\left(\frac{R_i}{h}\right). \tag{1.11}$$

In finite samples this estimator can produce discontinuous distribution functions that pose problems for inversion. To ensure the estimated function is continuous and to optimize the tradeoff between bias and variance, I follow Yu and Jones's (1998) suggestion to smooth "in the y-direction", as well, so that the local linear estimator of $\lim_{r\to 0^-} F_{Y|D=0,R=r}(y)$ would be $\hat{F}^{YJ}_{Y|D=0,R=0^-}(y)$, where:

$$\left(\hat{F}^{YJ}_{Y|D=0,R=0^-}(y), \hat{b}_{YJ}(y)\right) = \arg\min_{a,b} \sum_{i:D=0,Z=0} \left[\Omega\left(\frac{y-Y_i}{h_2}\right) - a - bR_i\right]^2 K\left(\frac{R_i}{h_1}\right),$$

where $\Omega(\cdot)$ is the distribution function associated with a kernel density function, W. This is essentially a conditional *distribution* version of Fan, Yao, and Tong's (1996) "double-kernel" conditional *density* estimator. This double-smoothed estimator has the following

[6]More generally, higher-order polynomials may also certainly be contemplated, but the choice of how many terms to include introduces another smoothing parameter, of which there are already several in the current approach. For simplicity, and also because it works well in practice, I focus on the local linear case.

closed form:

$$\hat{F}^{YJ}_{Y|D=0,R=0^-}(y) = \frac{1}{\sum_{j:D=0,Z=0} w_j(h_1)} \sum_{j:D=0,Z=0} w_j(h_1)\,\Omega\left(\frac{y-Y_j}{h_2}\right), \tag{1.12}$$

where the weighting function associated with local linear fitting is given by:

$$w_j(h_1) = K\left(\frac{R_j}{h_1}\right)[S_{n,2} - R_j S_{n,1}],$$

with

$$S_{n,l} = \sum_{i:D=0,Z=0} K\left(\frac{R_i}{h_1}\right) R_i^l, \quad l = 1, 2.$$

The bandwidths used for smoothing in the x-direction and y-direction are, respectively, h_1 and h_2. Yu and Jones (1998) give operational rules of thumb for choosing these bandwidths:

$$\begin{aligned} h_{1,p} &= h_{\text{mean}}\left\{p(1-p)/\phi\left(\Phi^{-1}(p)\right)^2\right\}^{1/5} \\ h_{2,p} &= \max\left(\frac{h^5_{1,1/2}}{h^3_1}, \frac{h_{1,p}}{10}\right) \text{ if } h_{1,1/2} < 1 \\ &\quad \text{and } \frac{h^4_{1,1/2}}{h^3_{1,p}} \text{ otherwise,} \end{aligned}$$

where p is the quantile index being estimated, ϕ and Φ are the standard normal density and cdf, respectively, and h_{mean} is a suitable bandwidth for estimating a conditional mean at a boundary.[7] Occasionally the quantile estimator (1.6) may not be monotone everywhere in τ (the quantile crossing problem) in which case the rearrangement technique of Chernozhukov, Fernández-Val, and Galichon (2009) may be applied to obtain a monotonic estimate of the quantile function. Chernozhukov, Fernández-Val, and Galichon (2009) show this improves the finite sample properties of the estimator, and has the same first order asymptotic distribution.

The conditional expectations in (1.10) are precisely the same quantities as those in the denominator of Hahn, Todd, and van der Klaauw's (2001) local Wald estimator, and can be estimated by local linear regression as they suggest. The bandwidth may be chosen by the plug in method to minimize the mean squared error as I describe in detail in the Appendix.

[7] I use a plug-in estimator based on an adaptation of Ruppert, Sheather, and Wand (1995) to the regression discontinuity setting. See the appendix for details.

1.4.1 Special Case: Sharp and Semi-fuzzy Designs

The identification results and estimation procedure described above assumes a fuzzy RD design, but they apply equally well to the special cases of sharp and semi-fuzzy designs, where treatment status may be constant on one or both sides of the threshold. Identification is more straightforward, since the monotonicity condition, Assumption 3, is automatically satisfied in the sharp and semi-fuzzy designs. In the sharp design case, the procedure still consists of inverting estimated distribution functions on each side of the threshold, but the expressions for the distribution functions to be inverted simplify to $F_{Y_1|C,R=0}(y) = \lim_{r \to 0^+} F_{Y|D=1,R=r}(y)$ and $F_{Y_0|C,R=0}(y) = \lim_{r \to 0^-} F_{Y|D=0,R=r}(y)$. The recommended procedure for estimating these conditional distributions still holds.

1.4.2 Including Additional Covariates

Empirical analysis often includes additional covariates for any of three reasons: to establish identification, to explore how the parameter of interest varies in subgroups, or to increase precision. Perhaps the main advantage of the canonical RD design as described above is identification does not require controlling for additional covariates. One situation where control might seem necessary arises when covariates appear different on either side of the discontinuity, even very near the threshold. In this case one might suspect the covariates are actually outcomes in and of themselves, or that the research design is flawed (e.g., location on either side of the threshold does not approximate quasi-experimental variation because, say, it is a choice). In such situations controlling for covariates may either be harmful or bespeak deeper problems with the research design. Thus including additional covariates for the sake of identification does not seem desirable.

The second common motivation for including covariates—exploring how effects vary in subgroups—can be incorporated in a straightforward manner. In the case of discrete covariates, all the results in the paper hold equally well conditional on a covariate X taking on a specific value, say, x. Naturally, the interpretation (and perhaps relevance) of the estimand changes when conditioning on covariates. Specific values of a continuous covariate may also be conditioned upon, but this would require another smoothing parameter, and a curse of dimensionality would apply, slowing the rate of convergence and reducing power. Conditioning on specific values of a continuous covariate therefore does not appear especially useful in this context.

The third motivation for including covariates—to increase efficiency—is potentially applicable in this context. Froelich (2007) shows for the case of mean estimation in the RD design how conditioning on covariates and integrating them out increases efficiency. Because all covariates but the running variable are integrated out, the curse of dimensionality is avoided and the rate of convergence is not affected. A necessary condition for including covariates to increase efficiency is that the covariate be a predictor of the outcome variable conditional on being at the threshold, and that its distribution conditional on the running variable be continuous at the threshold. These conditions are easily tested, and if they hold, one may directly apply Froelich's (2007) method, but replacing his estimand with (1.2) and (1.3), and inverting to obtain the local quantile treatment effect.

1.5 Asymptotic Distribution Theory

In this section I derive the limiting distribution for the local linear quantile treatment effects estimator, (1.5), obtained via inverting estimated distribution functions. First I will define some notation that will be useful for the theorem, as well as additional regularity assumptions. The notation and regularity assumptions I employ are analogous to those used by Hahn, Todd, and van der Klaauw (1999). Define $m_d(y,r) = F_{Y|D=d,R=r}(y)$ for $d \in \{0,1\}$ and $p(r) = E[D|R=r]$. Define the limits $m_d^+(y,r) = \lim_{e\to 0^+} m_d(y, r+e)$, $m_d^-(y,r) = \lim_{e\to 0^+} m_d(y, r-e)$, $p^+(r) = \lim_{e\to 0^+} p(r+e)$, $p^-(r) = \lim_{e\to 0^+} p(r-e)$. Additionally, define

$$\sigma_d^{2+}(y) = m_d^+(y,0)\left(1-m_d^+(y,0)\right), \quad \sigma_d^{2-}(y) = m_d^-(y,0)\left(1-m_d^-(y,0)\right).$$

Finally, define the following constant:

$$\omega^+ = \frac{\int_0^\infty (s_2 - s_1 u)^2 K(u)^2 du}{f_R(0)\gamma_{h_1}\cdot\left[s_2 s_0 - (s_1)^2\right]^2}, \tag{1.13}$$

where $s_l = \int_0^\infty K(u) u^l du$. Define ω^- similarly, but now with the integral in the limits of integration over $(-\infty, 0)$.

In addition to assumptions 1-3, I make the following assumptions in the derivation of the limiting distribution:

A1 For $d \in \{0,1\}$, the collections of functions $\{m_d(y,r)\}_y$, and $\{p(y,r)\}_y$ (viewed as functions of r indexed by y) are twice continuously differentiable. Uniformly in y, there exists some $M > 0$ such that $\left|m_d'^{+}(y,r)\right|$, $\left|m_d''^{+}(y,r)\right|$, and $|p'^{+}(r)|$, $|p''^{+}(r)|$ are uniformly bounded on $r \in (0, M]$. Similarly, uniformly in y, $\left|m_d'^{-}(y,r)\right|$, $\left|m_d''^{-}(y,r)\right|$, and $|p'^{-}(r)|$, $|p''^{-}(r)|$ are uniformly bounded on $r \in [-M, 0)$.

A2 The limits $m_d^{+}(y,0)$, $m_d^{-}(y,0)$, $m_d'^{+}(y,0)$, $m_d'^{-}(y,0)$, $m_d''^{+}(y,0)$, $m_d''^{-}(y,0)$, $p^{+}(0)$, $p^{-}(0)$, $p'^{+}(0)$, $p'^{-}(0)$, $p''^{+}(0)$, $p''^{-}(0)$ exist, for $d \in \{0,1\}$, and uniformly in y.

A3 The density of R, $f_R(r)$, and the conditional densities $f_{Y|D=d,R=r}(y)$ for $d \in \{0,1\}$ are continuous and bounded (uniformly in y) near 0. It is also bounded away from zero near 0. R and Y have bounded support.

A4 $K(\cdot), W(\cdot)$ are Borel measurable, bounded, continuous, symmetric, nonnegative-valued with compact support.

A5 The limits $\sigma_d^{2+}(y)$ and $\sigma_d^{2-}(y)$ exist, uniformly in y.

A6 $\lim\limits_{r\to 0} E\left[\left|1(Y \leq y) - m_d(y,r)\right|^3 \middle| D = d, R = r\right]$ exists, uniformly in y.

A7 The bandwidth sequences satisfy $h_1 = \gamma_{h_1} \cdot n^{-b}$ and $h_2 = \gamma_{h_2} h_1$ for some γ_{h_1} and γ_{h_2}, with $\frac{1}{5} < b < 1$.

Assumptions A1, A2, and A5 ensure the underlying conditional distributions of potential outcomes are sufficiently smooth to be well-approximated by local linear functions, and well-behaved at the discontinuity. A3 ensures that the distributions of the running variable and dependent variable are well-behaved near the threshold. The assumption of bounded support for R is without loss of generality, since the sequence of bandwidths ignores observations far from the discontinuity threshold. A4 is common in the kernel estimation literature and ensures integrals involving the kernel functions converge. A6 allows us to apply Lyapounov to obtain a functional central limit theorem. A7 specifies the sequence of bandwidths such that the nonparametric estimates converge but are asymptotically unbiased. Asymptotic normality of the quantile treatment effect estimator, (1.5), follows from the weak convergence as a process of conditional distribution function estimators, such as (1.11), and the functional delta method (van der Vaart, 1998). The following lemma establishes the weak convergence as a process of the local linear conditional distribution function estimator, (1.11).

Lemma 2 *Under assumptions 1-3 and A1-A7 the sequence*

$$\left\{\sqrt{nh_n}\left(\hat{F}_{Y|D=0,R=0^-}(y) - F_{Y|D=0,R=0^-}(y)\right) : y \in \mathbb{R}\right\}$$

is asymptotically tight in $\ell^\infty(\mathbb{R})$ *and converges in distribution to a Gaussian process.*

This result may be of interest in its own right, as it would be the basis of deriving the the limiting distribution of any parameter that is a function of the counterfactual distributions, including distribution treatment effects, counterfactual densities, and comparisons of stochastic dominance. Having established that the local linear conditional distribution function estimators converge as processes, we can apply a functional delta method to derive the limiting distribution of the quantile treatment effect estimator in the following theorem.

Theorem 3 *LQTE Asymptotic Distribution.*

Let c be the vector $\begin{pmatrix} 1 & -1 \end{pmatrix}'$. *Then under assumptions 1-3 and A1-A7 and for* $\tau \in (0,1)$*, the local linear quantile treatment effects estimator, (1.5) is asymptotically normally distributed with a limiting distribution given by the following:*

$$n^{\frac{1-b}{2}}\left(\hat{\delta}_{LQTE}(\tau) - \delta_{LQTE}(\tau)\right) \xrightarrow{d} N\left[0, c' J_{F_C} J_P \Sigma_P J_P' J_{F_C}' c\right], \tag{1.14}$$

where $1/5 < b < 1$*, as in A7, the matrices* J_{F_C}*,* J_P*, and* Σ_P *are as follows:*

$$J_{F_C} = \begin{pmatrix} \left(f_{Y_1|C,R=0}\left(Q_{Y_1|C,R=0}(\tau)\right)\right)^{-1} & 0 \\ 0 & \left(f_{Y_0|C,R=0}\left(Q_{Y_0|C,R=0}(\tau)\right)\right)^{-1} \end{pmatrix}, \tag{1.15}$$

$$J_P = \frac{1}{p^+(0) - p^-(0)} \begin{pmatrix} p^+(0) & 0 \\ -p^-(0) & 0 \\ 0 & 1 - p^-(0) \\ 0 & -\left(1 - p^+(0)\right) \\ D_1^+\left(Q_{Y_1|C,R=0}(\tau)\right) & D_0^+\left(Q_{Y_0|C,R=0}(\tau)\right) \\ -D_1^-\left(Q_{Y_1|C,R=0}(\tau)\right) & -D_0^-\left(Q_{Y_0|C,R=0}(\tau)\right) \end{pmatrix}', \tag{1.16}$$

where

$$D_d^{+/-}(y) = \lim_{r \to 0^{+/-}} F_{Y|D=d,R=r}(y) - F_{Y_d|C,R=0}(y),$$

and Σ_P *is a diagonal matrix with diagonal elements given by* $\omega^+\sigma_1^{2+}(y)$, $\omega^-\sigma_1^{2-}(y)$, $\omega^-\sigma_0^{2-}(y)$, $\omega^+\sigma_0^{2+}(y)$, $\omega^+p^+(0)(1-p^+(0))$, *and* $\omega^-p^-(0)(1-p^-(0))$, *respectively.*

1.6 Inference

In this section I discuss hypothesis testing and construction of confidence intervals for the local linear quantile treatment effects estimator, (1.5). One method is to consistently estimate the asymptotic variance derived in Section 1.5, and apply the asymptotic normality result to test hypotheses or construct confidence intervals. A second method is to use the nonparametric bootstrap.

Consistent estimation of the asymptotic variance requires consistent estimates of the components the matrices J_{F_C}, J_P, and Σ_P given in Theorem 3. Each element in these matrices is estimated as an intermediate quantity in the estimation procedure suggested in Section 1.4, with the exceptions of $f_{Y_1|C,R=0}(Q_{Y_1|C,R=0}(\tau))$ and $f_{Y_0|C,R=0}(Q_{Y_0|C,R=0}(\tau))$ in J_{F_C}, and $f_R(0)$, $\sigma_d^{2+}(0)$, and $\sigma_d^{2-}(0)$ in Σ_P. The conditional densities $f_{Y_1|C,R=0}(Q_{Y_1|C,R=0}(\tau))$ and $f_{Y_0|C,R=0}(Q_{Y_0|C,R=0}(\tau))$ can be written analogously to the conditional distribution functions (1.8) and (1.9), but substituting in the corresponding density functions for the distribution functions there:

$$\begin{aligned}
f_{Y_1|C,R=0}(y) &= \frac{\lim\limits_{r\to0^+} f_{Y|D=1,R=r}(y)\lim\limits_{r\to0^+} E[D|R=r] - \lim\limits_{r\to0^-} f_{Y|D=1,R=r}(y)\lim\limits_{r\to0^-} E[D|R=r]}{\lim\limits_{r\to0^+} E[D|R=r] - \lim\limits_{r\to0^-} E[D|R=r]} \\
f_{Y_0|C,R=0}(y) &= \frac{\lim\limits_{r\to0^+} f_{Y|D=0,R=r}(y)\lim\limits_{r\to0^+} E[1-D|R=r] - \lim\limits_{r\to0^-} f_{Y|D=0,R=r}(y)\lim\limits_{r\to0^-} E[1-D|R=r]}{\lim\limits_{r\to0^+} E[1-D|R=r] - \lim\limits_{r\to0^-} E[1-D|R=r]}.
\end{aligned}$$

The densities $\lim_{r\to0^+} f_{Y|D=1,R=r}(y)$, $\lim_{r\to0^-} f_{Y|D=1,R=r}(y)$, $\lim_{r\to0^-} f_{Y|D=0,R=r}(y)$, and $\lim_{r\to0^+} f_{Y|D=0,R=r}(y)$ in these formulae, as well as $f_R(0)$ can be consistently estimated using Fan, Yao, and Tong's (1996) local linear conditional density estimator. The conditional variances $\sigma_d^{2+}(0)$ and $\sigma_d^{2-}(0)$ can be estimated using local linear regression (Fan, 1992). After constructing consistent estimates of the asymptotic variance from these intermediate consistent estimates, hypothesis tests or confidence intervals based on the normal distribution will be consistent.

The second method for performing inference is the nonparametric bootstrap. Validity of the bootstrap in this setting depends on (a) the ability of the bootstrap to consistently estimate the distributions of the component nonparametric conditional mean estimators of the form (1.11), and (b) the applicability of a bootstrap delta method for valid inference on the final estimand. Shao and Tu (1995) and Horowitz (2001a) show the bootstrap is valid for nonparametric kernel conditional mean estimators if the asymptotic bias is removed, as it will be if the bandwidth is chosen to satisfy the conditions of Theorem 3. The first condition for bootstrap validity therefore is likely to be satisfied here.[8] The second condition, a bootstrap delta method, would follow from an extension of Theorem 3.9.11 in van der Vaart and Wellner (1996) to nonparametric estimators. This theorem requires that the functional defining the final estimator be Hadamard differentiable, and that the component estimators converge as processes, both of which are shown in the proof of Theorem 3. For the purposes of this paper I conjecture that the result extends, and I accordingly use the nonparametric bootstrap for inference in the application. Advantages of the bootstrap over estimating the asymptotic variance include that accomodating features of the data such as clustering is straightforward for the bootstrap, as well as availability of software implementations of the bootstrap.

1.7 Monte Carlo results

To illustrate the practical performance of the various procedures for estimating quantile treatment effects in the RD design, in this section I present the results of Monte Carlo simulations. The primitive of the model underlying the simulations is the joint distribution of (Y_0, Y_1, D, R), which I specify as follows: $R \sim N\left(0, \sigma_R^2\right)$, $Y_0 = R + \varepsilon_0$, $Y_1 = Y_0 - \varepsilon_1$, $D = 1\left(Y_1 - Y_0 + \gamma 1\left(R > 0\right) \geq \varepsilon_D\right)$, and the disturbance terms $(\varepsilon_0, \varepsilon_1, \varepsilon_D)$ are jointly normal and independent with mean zero and variances σ_0^2, σ_1^2, and σ_D^2, respectively. This model could be interpreted as a simple Roy model of selection on gains where exceeding the threshold $R = 0$ reduces the gross cost of treatment, ε_D, by γ. It exhibits the key features of the RD design with heterogeneous treatment effects. Note that the average treatment effect (ATE) is zero. In this model, the complier group is $C = \{0 < \varepsilon_D + \varepsilon_1 \leq \gamma\}$, and for positive γ, the local average treatment effect (LATE) is negative.

[8]Shao and Tu (1995) and Horowitz (2001a) discuss locally constant estimators, but the arguments go through for local linear estimators as well. Consistent with this, Kim and Truong (1998) apply the bootstrap to local linear smoothers.

Two key parameters affecting the performance of the estimators are the sample size, N, and the change in the probability of treatment at the threshold, Δp. The greater the magnitude of Δp, the higher the precision of the estimator. In the simulation model, the parameter γ controls this change in probability:

$$\Delta p = \Phi\left(\frac{\gamma}{\sqrt{\sigma_1^2 + \sigma_D^2}}\right) - \Phi(0).$$

I illustrate the performance of the estimators for several scenarios which are broadly representative of actual empirical examples. For each scenario, I perform 500 repetitions with parameter values $\sigma_R = \sigma_0 = \sigma_1 = \sigma_D = 1$, with $\gamma = .5$ ("small Δp") or $\gamma = 3$ ("large Δp") and a sample size of either $N = 10,000$ ("small N") or $N = 100,000$ ("large N"). To call a sample size of 10,000 "small" of course reflects the demands on the data nonparametric estimators require in general. I use a uniform kernel with bandwidths chosen as described in the appendix.

The key comparison is between local linear estimation of the quantiles of the compliers' potential outcomes and the "local constant" approach, (1.4) that applies instrumental variables quantile treatment effects estimation using kernel weights to narrow in on the threshold, ignoring the running variable (Abadie LQTE[9]). Within the local linear approaches, I performed simulations for the three alternatives mentioned in Section 1.4: (1) direct quantile estimation by local linear check function minimization; (2) inverting local linear 2SLS estimates of the distribution functions; and (3) the preferred approach of inverting estimates of the distribution functions obtained by estimating the components separately (RD LQTE). In all scenarios the check function approach had virtually identical performance to the preferred approach, except it took longer to compute. Compared to the preferred approach, check function minimization required 55 percent more cpu time and more than 3.5 times as many iterations to converge (on average 19.8 iterations per quantile estimate versus 5.6 iterations for the preferred approach). The local linear 2SLS approach also performed nearly identically to the preferred approach, except in scenarios with smaller sample sizes and unbalanced numbers of treated units above and below the threshold ("large Δp"), where it performed slightly worse (that is, greater bias and variance). Accordingly, I present results

[9] Although I refer to the application of Abadie, Angrist, and Imbens's (2002) IV quantile treatment effects estimator to the RD setting using the shorthand "Abadie LQTE", I emphasize those authors did not propose applying their estimator to the RD design.

from the preferred technique as indicative of the local linear approaches, and I show results from the local linear 2SLS approach only when it differs perceptibly from the preferred approach.

The first simulation scenario, "large N, large Δp", represents the most favorable conditions for the estimation procedures. I set $\gamma = 3$, which implies a jump in the probability of treatment at the threshold of about 48 percent, and I use a sample size of 100,000. This change in the probability and the sample size are on the order of those found in several recent RD studies, including Matsudaira (2008) and Jacob and Lefgren (2004). In this and all scenarios, bandwidths for the preferred approach were chosen in each repetition according to the procedure outlined in Appendix B. Typical values for h_{mean} were in the range .15-.3. The bandwidths for the local constant approach were chosen to minimize the simulated mean squared error. While this choice is infeasible in practice, it gives an upper bound on how well a local constant approach might perform relative to local linear approaches.

Figure 1-1 shows the results of the simulation under this scenario for RD LQTE. The figure shows the average point estimate for each quantile index, as well as the pointwise (in the quantile index) 90 percent confidence interval spanned by the fifth and 95th percentile estimate. The figure shows that the bias is very small, despite the fact that the estimator consists of nonlinear functions of estimated quantities. The estimates also appear to be quite precise, although the simulations are not calibrated to any particular economic context to give the scale meaning. By way of comparision, Figure 1-2 illustrates the performance of applying the locally constant Abadie LQTE approach to the RD setting. The confidence intervals are somewhat wider, and the estimates are substantially biased. Thus in terms of bias and variance, the local linear RD LQTE approach appears to do strictly better, although for samples this large and discontinuities of this size the difference is not extremely large.

The second simulation scenario, "small N, large Δp" illustrates the implications for the performance of the estimators when the sample is smaller. The change in the probability of treatment remains at 48 percent, but I use a sample size of 10,000. This corresponds roughly to the empirical application in Section 1.8 based on Gormley, Gayer, Phillips, and Dawson (2005). Figure 1-3 shows the confidence intervals for RD LQTE are substantially wider than for the large N case, but the bias remains negligible. Figure 1-4 shows Abadie LQTE, on the other hand, has confidence intervals around 30 percent wider still than RD

LQTE, and the bias is also larger than for the large N case. This scenario represents a situation where the approach of separately estimating the components of the distribution functions (1.8) and (1.9) may have advantages over the local linear 2SLS approach, since there are few untreated units above the threshold and many below the threshold. The local linear 2SLS approach implicitly uses the same bandwidth above and below, while the preferred approach allows separate bandwidths to be chosen optimally. Figure 1-5 shows that the implicit single bandwidth slightly worsens performance. Compared to the preferred approach, the local linear 2SLS approach exhibits slightly more bias and variance, although the difference is small.

Finally, the third scenario, "large N, small Δp" preserves the large sample size of 100,000, but sets $\gamma = .5$, implying a change in the probability of treatment at the threshold of about 14 percent. Figure 1-6 shows the results from this scenario for RD LQTE. Despite the large sample size in this scenario, the smaller Δp results in the widest confidence intervals for any of the simulation scenarios I consider. The bias, however, remains negligible, even for a much smaller Δp. As Figure 1-7 shows, the confidence intervals for the applying Abadie LQTE to RD are also widest for this scenario, and the bias is the greatest as well. Thus even for relatively large sample sizes, a small jump at the threshold in the probability of treatment can result in significant bias in the local constant approach.

In summary, the simulation exercises highlight that the proposed estimation procedure requires relatively large samples, a large discontinuity in the probability of treatment, or, ideally, both, to attain good precision. In all cases, however, the bias was minimal. The simulations also highlight that local linear approaches perform unambiguously better than an approach ignoring the effects of R within a window around the discontinuity.

1.8 Application: Effects of Universal Pre-K

In this section I apply the RD quantile treatment effects procedure to an example from the literature which will both illustrate how the procedure might be applied to real-life questions, as well as point out some challenges faced by nonparametric estimation of distributional effects.

Interventions designed to improve educational performance are one setting in which distributional effects may be important to policy makers. One such policy that specifically targets the lower end of the distribution is the introduction of universal pre-K programs. Gormley, Gayer, Phillips, and Dawson (2005) use a regression discontinuity design to analyze an Oklahoma universal pre-K program, and find significant positive effects on average test scores measuring cognitive development along a variety of dimensions. By conditioning on various socio-economic status indicators, they find indirect suggestive evidence that the program also has positive effects on the lower end of the distribution. The quantile treatment effects estimator developed in this paper allows direct investigation of the effect of the policy on the lower end of the distribution.

Oklahoma introduced a universal pre-K program for four-year-olds in 1998, and by 2002-2003 (the period I analyze) 91 percent of the state's school districts were participating, including Tulsa Public Schools (TPS), the largest district in the state, and the district from which my sample is drawn.

A child's participation in the pre-K program is voluntary (on the part of the parents), but is subject to a birthday cutoff eligibility rule. Children who had turned four years old by September 1, 2002 were eligible for the program, while younger children were not. Figure 1-8 shows the discontinuity in probability of treatment that the eligibility rule induced. Because the participation among children who missed the cutoff is essentially nil, local treatment effects in this setting correspond to the effect of treatment on the treated.

At the start of the 2003-2004 school year, all incoming kindergartners and TPS pre-K participants were given the Woodcock-Johnson Achievement Test, a nationally normed test that has been widely used in studies of early education. Treated students are those who participated in a TPS pre-K program the previous year.

I use a sample of 4,710 incoming TPS kindergartners and pre-K participants. The dataset includes exact date of birth, an indicator for participation in TPS pre-K the previous year (the treatment variable), and scores on the three Woodcock-Johnson subtests: Letter-word, Spelling, and Applied Problems.

Using a uniform kernel and bandwidths chosen by the plug-in method described in the Appendix, I estimated the local treatment effects of attending the Pre-K program on the three subtests of the Woodcock-Johnson test. The estimated optimal bandwidths were on the order of 30 to 40 days. The estimated local quantile treatment effects of TPS pre-K programs on scores on the three subtests of the Woodcock-Johnson tests are plotted in figure 1-9. The figure shows the point estimates for each quantile index, as well as pointwise (in the quantile index) 90 percent confidence intervals from bootstrap simulations.

Panel A shows a relatively precisely estimated two to four point (around 80 percent of a standard deviation) effect on the lower end of the distribution of the Letter-Word Identification score, with the effects in the middle of the distribution somewhat larger than at the lowest end. Effects on the upper end of the distribution are less precisely estimated, but the point estimates decline at the upper end, and the effect ceases to be significantly different from zero. We cannot rule out, however, that the effects on the upper end of the distribution are as large (or larger) than those at lower points in the distribution. Effects across the distribution of the Spelling score are plotted in panel B, and are similar to the results for the Letter-Word Identification scores. Panel C shows the effects of pre-K participation on the Applied Problems score. The point estimates are largest and most precisely estimated for the bottom end of the distribution, and similarly to the other two subtests, point estimates of the effects for the top of the distribution are smaller and less precise. For reference, the local average treatment effects[10] (LATE) are estimated to be 3.66 for Letter-Word Identification, 1.93 for Spelling, and 3.44 for applied problems. An application of the local constant, Abadie-weighted approach (not reported) yields a similar pattern across quantiles, but with point estimates higher by about two tenths of a point, which is consistent with the bias observed in the simulations in Section 1.7. The alternative local linear 2SLS approach to estimating the distribution functions yielded results (not reported) very similar to those reported below for the preferred approach, but with somewhat wider confidence intervals, consistent with the comparison made in the simulations.

The estimation results imply that universal pre-K in Oklahoma succeeded in significantly raising the lower end of the distribution of test scores, especially for the Applied Problems subtest. These results are consistent with the Gormley, Gayer, Phillips, and Dawson's (2005) findings that point estimates of average effects were larger for children from potentially disadvantaged socio-economic groups. These results are subject to the caveat that they measure the net effect of participating in a TPS pre-K program versus alternatives parents might have chosen in absence of the program. The alternatives may have been different for children at different points in the distribution, and thus we cannot draw conclusions about

[10]LATE was computed as in Hahn, Todd, and van der Klaauw (2001): $\hat{\delta}_{LATE} = \left\{\hat{E}\left[Y|R=0^+\right] - \hat{E}\left[Y|R=0^-\right]\right\} / \left\{\hat{E}\left[D|R=0^+\right] - \hat{E}\left[D|R=0^-\right]\right\}$.

the gross impact of universal pre-K programs on the distribution of outcomes. An additional caveat is that these results reflect the short-term effect. It's possible that children who did not participate may catch up over time, although evidence from the Perry Preschool Study (Schweinhart, Barnes, Weikart, Barnett, and Epstein, 1993; Anderson, 2008) suggests there may be significant long term impacts of pre-K programs.

This application illustrated the ability of the estimation procedure to evaluate distributional policy effects, but it also highlighted some challenges involved with nonparametric estimation in general, and especially with nonparametric estimators of distributional effects. Meaningful inference requires large samples. The imprecision of the estimates for the upper end of the distribution reflects the modest sample size compounded by the skewness of the distribution of test scores, with much lower densities above the median.

1.9 Conclusion

In this paper I have introduced a new approach to estimating local quantile treatment effects in an RD design, showing consistency and asymptotic normality. The estimator is the horizontal difference between the marginal distributions of the potential outcomes for compliers, which are estimated via local linear quantile regression techniques. In contrast to other possible approaches to estimating distribution effects in an RD context, the procedure I have developed here relies only on the LATE assumptions, and avoids the significant finite sample bias that other "local constant" approaches suffer from, including the Abadie quantile regression kernel-weighted at the threshold. Monte Carlo simulations confirm that the bias of the approach I suggest is minimal compared to other approaches. The simulations also show that the estimation procedure performs best when the research design involves a large sample, a large discontinuity in the probability of treatment, or both.

An application of the procedure to estimating the effects of an Oklahoma universal pre-K program across the distribution of test scores shows that the lower end of the distribution is significantly raised, while estimates at the top of the distribution are smaller and less precise. Other possibilities for applying the methodology are numerous, and include the study of remedial education programs by Jacob and Lefgren (2004) and Matsudaira (2008), the study of the UI Worker Profiling and Reemployment Services program by Black, Smith, Berger, and Noel (2003), and the effect of unions on wages by DiNardo and Lee (2004). I leave the application of the RD quantile treatment effects estimation to these questions and others to future research.

Software implementations of the procedure developed in this paper are available at http://econ-www.mit.edu/grad/frandsen/software .

References

ABADIE, ALBERTO (2002): "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models," *Journal of the American Statistical Association*, 97, 284–292.

ABADIE, ALBERTO, JOSHUA D. ANGRIST, AND GUIDO IMBENS (2002): "Instrumental Variables Estimates of the Effect of Subsidized Training on the Quantiles of Trainee Earnings," *Econometrica*, 70, 91–117.

ANDERSON, MICHAEL (2008): "Multiple Inference and Gender Differences in the Effect of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, forthcoming.

ANGRIST, JOSHUA D., AND VICTOR LAVY (1999): "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *The Quarterly Journal of Economics*, 114, 533–575.

ATKINSON, ANTHONY B. (1970): "On the Measurement of Inequality," *Journal of Economic Theory*, 2, 244–263.

BAKER, MICHAEL, SERGIO FIRPO, AND KEVIN MILLIGAN (2005): "Identification and Estimation of Distributional Effects under a Regression Discontinuity Design with Application to Changes in Maternity Leave Mandates," Unpublished manuscript.

BLACK, DAN, JOSE GALDO, AND JEFFREY SMITH (2007): "Evaluating the Bias of the Regression Discontinuity Design Using Experimental Data," Mimeo, University of Chicago.

BLACK, DAN A., JEFFREY A. SMITH, MARK C. BERGER, AND BRETT J. NOEL (2003): "Is the Threat of Reemployment Services More Effective than the Services Themselves? Evidence from Random Assignment in the UI System," *The American Economic Review*, 93, 1313–1327.

BUDDELMEYER, H., AND E. SKOUFIAS (2003): "An Evaluation of the Performance of Regression Discontinuity Design on PROGRESA," IZA, Bonn, Germany.

CHERNOZHUKOV, VICTOR, IVÁN FERNÁNDEZ-VAL, AND ALFRED GALICHON (2009): "Quantile and Probability Curves Without Crossing," MIT Working Paper.

CHERNOZHUKOV, VICTOR, AND CHRISTIAN HANSEN (2005): "An IV Model of Quantile Treatment Effects," *Econometrica*, 73, 245–261.

COOK, THOMAS D., AND VIVIAN C. WONG (2008): "Empirical Tests of the Validity of the Regression-Discontinuity Design," *Annales d'Economie et de Statistique*, forthcoming.

DINARDO, JOHN, AND DAVID S. LEE (2004): "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 119, 1383–1441.

FAN, JIANQING (1992): "Design-adaptive Nonparametric Regression," *Journal of the American Statistical Association*, 87, 998–1004.

FAN, JIANQING, QIWEI YAO, AND HOWELL TONG (1996): "Estimation of Conditional Densities and Sensitivity Measures in Nonlinear Dynamical Systems," *Biometrika*, 83, 189–206.

FRANDSEN, BRIGHAM R. (2008a): "Long term effects of military service on the distribution of earnings," Mimeo, MIT.

——— (2008b): "Nonparametric Estimation of Local Quantile Treatment Effects in the Regression Discontinuity Design," Available at SSRN: http://ssrn.com/abstract=1216062.

FROELICH, MARKUS (2007): "Regression Discontinuity Design with Covariates," IZA Discussion Paper No. 3024.

FROELICH, MARKUS, AND BLAISE MELLY (2008): "Quantile Treatment Effects in the Regression Discontinuity Design," IZA Discussion Paper 3638, Institute for the Study of Labor (IZA).

GORMLEY, WILLIAM T. GORMLEY, JR., TED GAYER, DEBORAH PHILLIPS, AND BRITTANY DAWSON (2005): "The Effects of Universal Pre-K on Cognitive Development," *Developmental Psychology*, 41, 872–884.

GUITERAS, RAYMOND (2008): "Estimating Quantile Treatment Effects in a Regression Discontinuity Design," Mimeo, MIT.

HAHN, JINYONG, PETRA TODD, AND WILBERT VAN DER KLAAUW (1999): "Evaluating the Effect of an Antidiscrimination Law Using a Regression-Discontinuity Design," Working Paper 7131, National Bureau of Economic Research.

——— (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201–209.

HALL, P., S. J. SHEATHER, M. C. JONES, AND J. S. MARRON (1991): "On Optimal Data-Based Bandwidth Selection in Kernel Density Estimation," *Biometrika*, 78, 521–530.

HOROWITZ, JOEL L. (2001a): "The Bootstrap," in *Handbook of Econometrics*, ed. by James J. Heckman, and Edward E. Leamer, vol. 5, pp. 3159–3228. Elsevier Science, Amsterdam.

——— (2001b): "Nonparametric Estimation of a Generalized Additive Model with an Unknown Link Function," *Econometrica*, 69, 499–513.

IMBENS, GUIDO W., AND JOSHUA D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475.

IMBENS, GUIDO W., AND THOMAS LEMIEUX (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615–635.

IMBENS, GUIDO W., AND DONALD B. RUBIN (1997): "Estimating Outcome Distributions for Compliers in Instrumental Variables Models," *The Review of Economic Studies*, 64, 555–574.

JACOB, BRIAN A., AND LARS LEFGREN (2004): "Remedial Education and Student Achievement: A Regression Discontinuity Analysis," *The Review of Economics and Statistics*, 86, 226–244.

JONES, M. C. (1993): "Simple Boundary Correction for Kernel Density Estimation," *Statistics and Computing*, 3, 135–146.

Kim, Haesook T., and Young K. Truong (1998): "Nonparametric Regression Estimates with Censored Data: Local Linear Smoothers and Their Applications," *Biometrics*, 54, 1434–1444.

Matsudaira, Jordan (2008): "Mandatory Summer School and Student Achievement," *Journal of Econometrics*, 142, 829–850.

Porter, Jack (2003): "Estimation in the Regression Discontinuity Model," unpublished paper, Department of Economics, Harvard University.

Ruppert, D., S. J. Sheather, and M. P. Wand (1995): "An Effective Bandwidth Selector for Local Least Squares Regression," *Journal of the American Statistical Association*, 90, 1257–1270.

Schweinhart, L., H. Barnes, D. Weikart, W. S. Barnett, and A. Epstein (1993): *Significant benefits: Vol. 10. The High/Scope Perry Preschool Study Through Age 27.* High/Scope Press, Ypsilanti, MI.

Shao, Jun, and Dongsheng Tu (1995): *The jackknife and bootstrap.* Springer Verlag, New York, NY, USA.

van der Vaart, A.W. (1998): *Asymptotic Statistics.* Cambridge University Press, New York.

van der Vaart, A.W., and Jon A. Wellner (1996): *Weak convergence and empirical processes.* Springer-Verlag, New York.

Yu, Keming, and M. C. Jones (1997): "A comparison of local constant and local linear regression quantile estimators," *Computational Statistics & Data Analysis*, 25, 159–166.

——— (1998): "Local Linear Quantile Regression," *Journal of the American Statistical Association*, 93, 228–237.

1.10 Figures

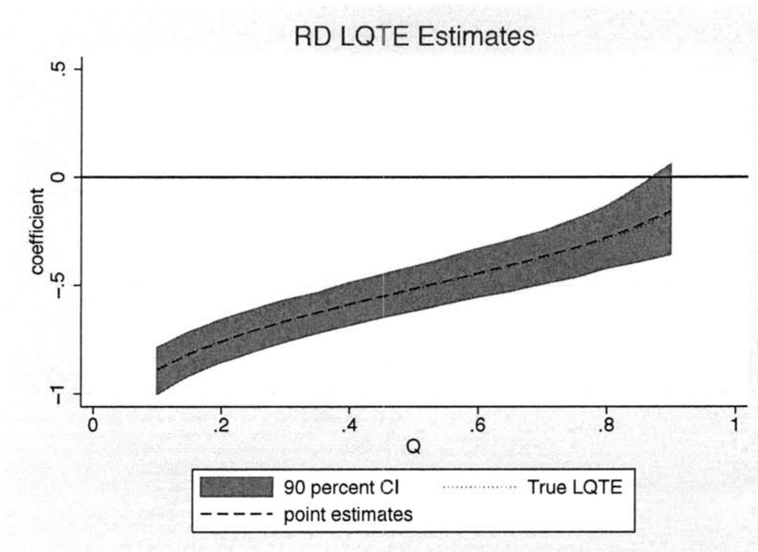

Figure 1-1: RD LQTE Monte Carlo Results: large N, large Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of RD LQTE with 500 repetitions, a sample size of 100,000, and a discontinuity in the probability of treatment at the threshold of 48 percent.

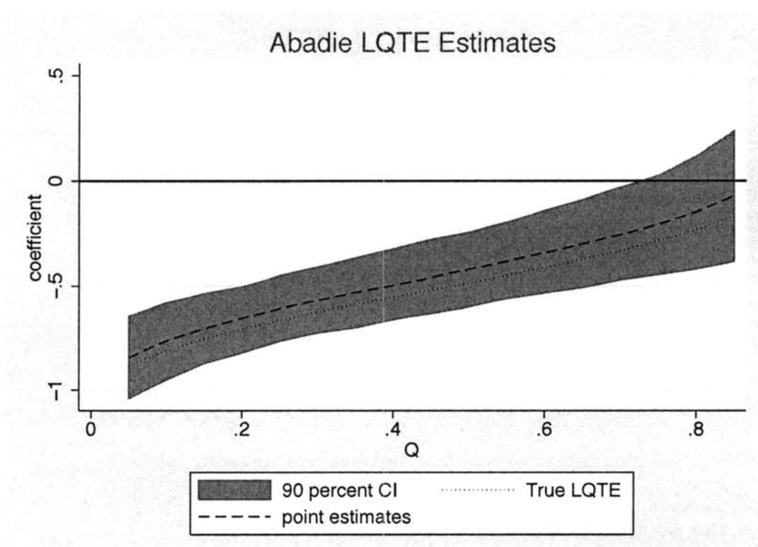

Figure 1-2: Abadie LQTE Monte Carlo Results: large N, large Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of kernel-weighted Abadie quantile regression with 500 repetitions, a sample size of 100,000, and a discontinuity in the probability of treatment at the threshold of 48 percent.

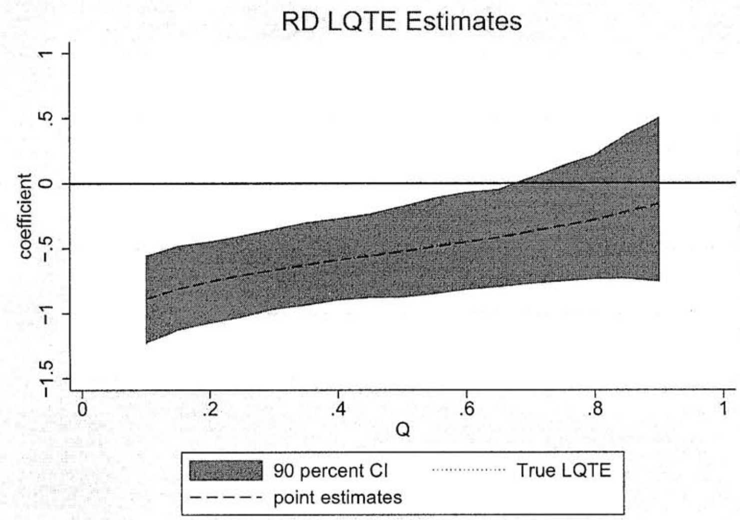

Figure 1-3: RD LQTE Monte Carlo Results: small N, large Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of RD LQTE with 500 repetitions, a sample size of 10,000, and a discontinuity in the probability of treatment at the threshold of 48 percent.

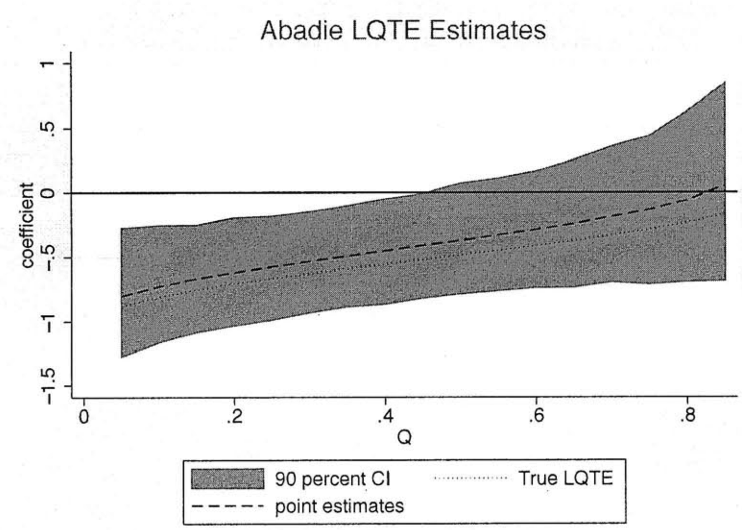

Figure 1-4: Abadie LQTE Monte Carlo Results: small N, large Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of kernel-weighted Abadie quantile regression with 500 repetitions, a sample size of 10,000, and a discontinuity in the probability of treatment at the threshold of 48 percent.

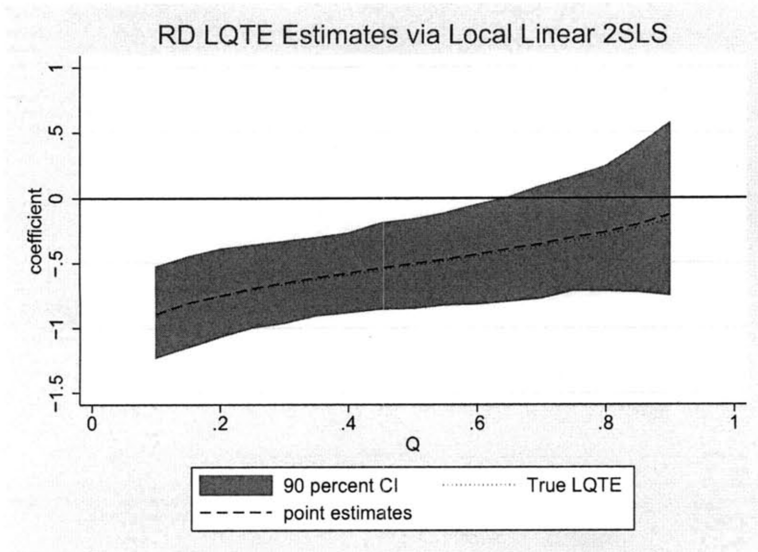

Figure 1-5: RD LQTE via local linear 2SLS Monte Carlo Results: small N, large Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of RD LQTE via local linear 2SLS with 500 repetitions, a sample size of 10,000, and a discontinuity in the probability of treatment at the threshold of 48 percent.

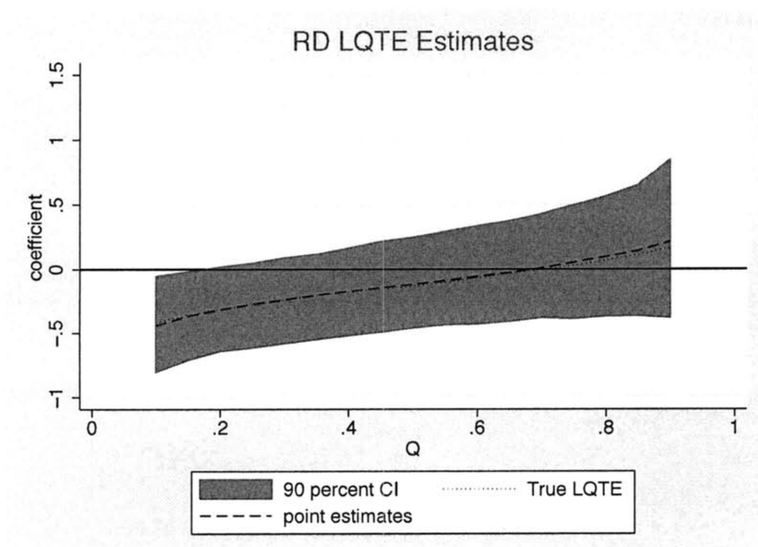

Figure 1-6: RD LQTE Monte Carlo Results: large N, small Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of RD LQTE with 500 repetitions, a sample size of 100,000, and a discontinuity in the probability of treatment at the threshold of 14 percent.

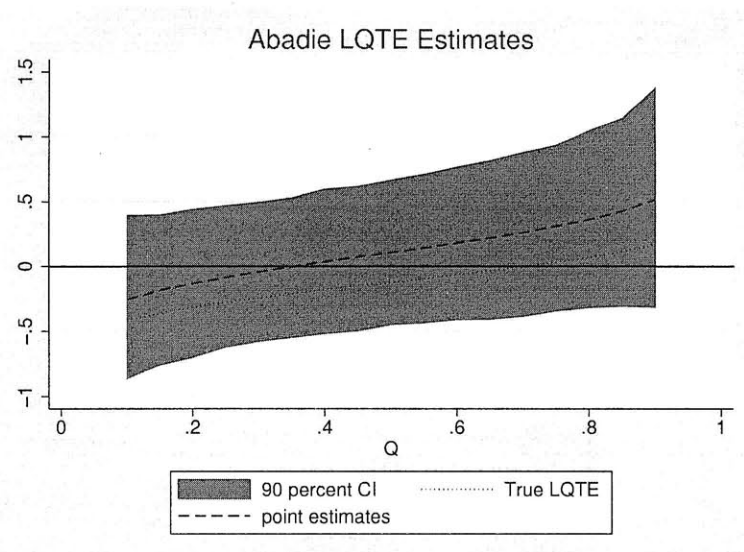

Figure 1-7: Abadie LQTE Monte Carlo Results: large N, small Δp. The figure shows point estimates and confidence intervals from a Monte Carlo simulation of kernel-weighted Abadie quantile regression with 500 repetitions, a sample size of 100,000, and a discontinuity in the probability of treatment at the threshold of 14 percent.

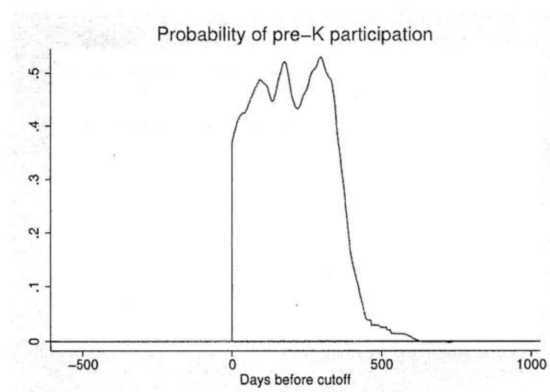

Figure 1-8: The figure plots the probability of attending TPS pre-K in 2002-2003 as a function of birthdate relative to cutoff

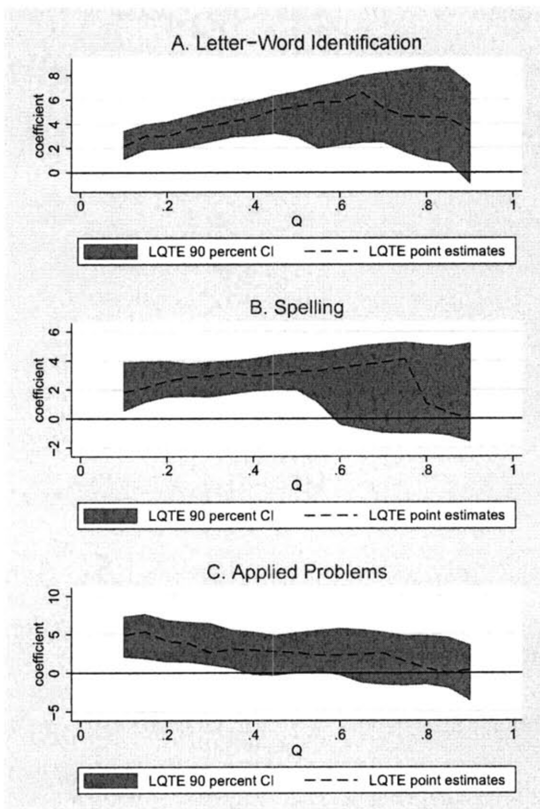

Figure 1-9: The figure plots point estimates and bootstrapped confidence intervals for the effect of TPS pre-K participation on the distribution of scores on three WJ subtests. A uniform kernel was used and bandwidths for each quantile index were chosen as described in the Appendix. The estimated optimal bandwidths were on the order of 30-40 days.

Appendix A: Proofs of lemmas and theorems

The main identification result, Theorem 1, follows from a special case of the following lemma. This lemma is an adaptation of Imbens and Rubin's (1997) and Abadie's (2002) results, but applied to the regression discontinuity context.

Lemma 4 *Expectations of Functions of Potential Outcomes for Compliers*

Let $h(\cdot)$ be any measurable function on the real line such that $E|h(Y)| < \infty$ and $E[h(Y_d)|D_1, D_0, R = r]$ is continuous and differentiable in r over an ε-neighborhood of zero for $d \in \{0,1\}$. Then under Assumptions 1-3, $E[h(Y_d)|C, R = 0]$, $\quad d \in \{0,1\}$ is identified from the joint distribution of (Y, D, R) as

$$E[h(Y_1)|C, R=0] = \frac{\lim_{r\to 0^+} E[h(Y)D|R=r] - \lim_{r\to 0^-} E[h(Y)D|R=r]}{\lim_{r\to 0^+} E[D|R=r] - \lim_{r\to 0^-} E[D|R=r]}$$

$$E[h(Y_0)|C, R=0] = \frac{\lim_{r\to 0^+} E[h(Y)(1-D)|R=r] - \lim_{r\to 0^-} E[h(Y)(1-D)|R=r]}{\lim_{r\to 0^+} E[1-D|R=r] - \lim_{r\to 0^-} E[1-D|R=r]}.$$

Proof. First we show the difference in the expected treatment status across the threshold is equal to the probability of being a complier in the neighborhood of the threshold:

$$\begin{aligned}
\lim_{r\to 0^+} E[D|R=r] - \lim_{r\to 0^-} E[D|R=r] &= \lim_{r\to 0^+} E[D_1|R=r] - \lim_{r\to 0^-} E[D_0|R=r] \\
&= E[D_1 - D_0|R=0] \\
&= \Pr(D_1 > D_0|R=0) \\
&= \Pr(C|R=0).
\end{aligned}$$

The first equality follows from the definition of potential treatment status. The second equality follows from the continuity of the conditional expectations of potential treatment in Assumption 2. The third equality follows from the fact that $D_1 - D_0$ equals one when $D_1 > D_0$ and zero when $D_1 = D_0$, and by the monotonicity condition, Assumption 3, $\Pr(D_1 < D_0|R=0) = 0$. The final equality follows from the definition of a complier. Next we show that the difference in the expectation of $h(Y)D$ across the threshold is equal to the expectation of $h(Y_1)$ for compliers, multiplied by the probability of being a complier, which establishes the first result:

$$\begin{aligned}
& \lim_{r \to 0^+} E\left[h(Y) D | R = r\right] - \lim_{r \to 0^-} E\left[h(Y) D | R = r\right] \\
= & \lim_{r \to 0^+} E\left[h(Y_1) D_1 | R = r\right] - \lim_{r \to 0^-} E\left[h(Y_1) D_0 | R = r\right] \\
= & E\left[h(Y_1)(D_1 - D_0) | R = 0\right] \\
= & E\left[h(Y_1) | D_1 > D_0, R = 0\right] \Pr(D_1 > D_0 | R = 0) \\
= & E\left[h(Y_1) | C, R = 0\right] \Pr(C | R = 0).
\end{aligned}$$

The first equality follows from the definitions of potential outcomes and potential treatment status. The second equality follows from the continuity of $E\left[h(Y_d) | D_1, D_0, R = r\right]$ and $E\left[D_d | R = r\right]$ (Assumption 2). The third equality follows from iterated expectations, and the final equality follows from the definition of compliers. The steps for $E\left[h(Y_0) | C, R = 0\right]$ are analogous. ■

Proof of Theorem 1

Proof. The proof follows from lemma 4. Let $h(Y_d) = 1(Y_d \leq y)$. Then by the smoothness condition, Assumption 2, $E\left[h(Y_d) | D_1, D_0, R\right] = F_{Y_d|D_1,D_0,R}(y|r)$ satisfies the smoothness hypothesis of the lemma, and by Assumptions 1-3, the remaining hypotheses of lemma 4 are satisfied, establishing that

$$F_{Y_1|C,R=0}(y) = \frac{\lim_{r \to 0^+} E\left[1(Y \leq y) D | R = r\right] - \lim_{r \to 0^-} E\left[1(Y \leq y) D | R = r\right]}{\lim_{r \to 0^+} E\left[D | R = r\right] - \lim_{r \to 0^-} E\left[D | R = r\right]} \tag{1.17}$$

$$F_{Y_0|C,R=0}(y) = \frac{\lim_{r \to 0^+} E\left[1(Y \leq y)(1 - D) | R = r\right] - \lim_{r \to 0^-} E\left[1(Y \leq y)(1 - D) | R = r\right]}{\lim_{r \to 0^+} E\left[1 - D | R = r\right] - \lim_{r \to 0^-} E\left[1 - D | R = r\right]} \tag{1.18}$$

By Assumption 2, $F^{-1}_{Y_1|C,R=0}(\tau)$ and $F^{-1}_{Y_0|C,R=0}(\tau)$ exist, which establishes the result of the theorem. ■

Proof of Lemma 2

Proof. The local linear conditional distribution function estimator, (1.11), can be written as a function of sample averages:

$$\begin{aligned} \hat{F}_{Y|D=0,R=0^-}(y) &= \frac{A_{n,2}B_{n,0}(y) - A_{n,1}B_{n,1}(y)}{A_{n,2}A_{n,0} - A_{n,1}^2}, \qquad (1.19) \\ A_{n,l} &\equiv \frac{1}{n}\sum_{i=1}^{n}\frac{1}{h_n}K\left(\frac{R_i}{h_n}\right)\left(\frac{R_i}{h_n}\right)^l \\ B_{n,l}(y) &\equiv \frac{1}{n}\sum_{j=1}^{n}\frac{1}{h_n}1(Y_j \le y)K\left(\frac{R_j}{h_n}\right)\left(\frac{R_j}{h_n}\right)^l. \end{aligned}$$

For simplicity, in this proof I consider a sample drawn from the joint distribution of (Y, D, R) conditional on $D = 0, R \le 0$, with a sample size of n, and I therefore omit explicit conditioning on $D = 0, R \le 0$ in sums. I write the bandwidth as h_n to emphasize dependence on the sample size. To show the weak convergence of $\hat{F}_{Y|D=0,R=0^-}(y)$ I establish that each of the terms $A_{n,0}, A_{n,1}, A_{n,2}, B_{n,0}(y), B_{n,1}(y)$ converge weakly as processes, and apply a functional delta method. I start by establishing the convergence as a process of

$$\sqrt{nh_n}\left(B_{n,l}(y) - E\left[B_{n,l}(y)\right]\right), \quad l = 0, 1, \qquad (1.20)$$

since the $A_{n,l}$ terms are trivial functions of y. Define a vector of random variables, X_i, and indexing set T:

$$\begin{aligned} X_i &= \begin{pmatrix} Y_i \\ R_i \end{pmatrix}, \\ T &= \mathbb{R}. \end{aligned}$$

Define the set of functions $\mathcal{F}_n = \{f_{n,t} : t \in T\}$, with:

$$f_{n,t}(X_i) = 1(Y_i \le t)\frac{1}{\sqrt{h_n}}K\left(\frac{R_i}{h_n}\right)\left(\frac{R_i}{h_n}\right)^l, \quad l = 0, 1.$$

Then the process (1.20) can be written:

$$n^{-\frac{1}{2}}\sum_{i=1}^{n}\left(f_{n,t}(X_i) - Pf_{n,t}\right) : t \in T,$$

which corresponds to van der Vaart and Wellner's (1996) setup for Theorem 2.11.22 for convergence of processes indexed by classes of functions changing with n. Letting P and P^* denote measure and outer measure, respectively, and $\rho(s,t)$ a pseudonorm on $\mathbb{R}$, the conditions needed for convergence are the following:

1. There exist envelope functions $F_n : |f_{n,t}(x)| \leq F_n(x) \quad \forall x, f, n$ which satisfy (a) $P^* F_n^2 = O(1)$, and (b) $P^* F_n^2 \{F_n > \eta\sqrt{n}\} \to 0$, for every $\eta > 0$

2. $\mathcal{F}_{n,\delta} = \{f_{n,s} - f_{n,t} : \rho(s,t) < \delta\}$ and $\mathcal{F}_{n,\delta}^2$ are P-measurable for every $\delta > 0$

3. $f_{n,t}$ satisfy:
$$\sup_{\rho(s,t)<\delta_n} P(f_{n,s} - f_{n,t})^2 \to 0, \qquad \text{for every } \delta_n \downarrow 0,$$

4. The uniform entropy condition on page 220 of van der Vaart and Wellner holds.

Start with the first condition (envelope functions). Define a set of envelope functions to be:
$$F_n = \left| \frac{1}{\sqrt{h_n}} K\left(\frac{R_j}{h_n}\right)\left(\frac{R_j}{h_n}\right)^l \right|, \quad l = 0, 1.$$

Clearly these are envelope functions for class $\mathcal{F}_n$. Under the measurability assumption, condition 1a can be written:

$$\begin{aligned} PF_n^2 &= \int \left(\frac{1}{\sqrt{h_n}} K\left(\frac{R_j}{h_n}\right)\left(\frac{R_j}{h_n}\right)^l \right)^2 dF_{R|R\leq 0, D=0}(r), \quad l = 0, 1 \\ &= \int \left(K(u)(u)^l\right)^2 f_R(h_n u)\, du, \quad l = 0, 1, \end{aligned}$$

making the change of variables $u = \frac{r}{h_n}$. Condition 1a then holds under our boundedness assumptions on R and $K(\cdot)$. Condition 1b holds trivially for $l = 0$ for bounded $K(\cdot)$. For $l = 1$, 1b is essentially the Lindeberg-Feller condition, and holds if, for example, R is bounded. Condition 2 is implied by our assumption that $K(\cdot)$ is measurable.

The quantity in condition 3 can be written:

$$
\begin{aligned}
& \sup_{\rho(s,t)<\delta_n} P\left(f_{n,s}-f_{n,t}\right)^2 \\
= & \sup_{\rho(s,t)<\delta_n} P\left(\left(1\left(Y_i \leq s\right)-1\left(Y_i \leq t\right)\right) \frac{1}{\sqrt{h_n}} K\left(\frac{R_i}{h_n}\right)\left(\frac{R_i}{h_n}\right)^l\right)^2 \\
= & \int_{r \leq 0}\left\{\sup_{\rho(s,t)<\delta_n} \int_y (1(y \leq s)-1(y \leq t))^2 d F_{Y \mid R=r, R \leq 0, D=0}(y)\right\} \\
& \times\left(\frac{1}{\sqrt{h_n}} K\left(\frac{r}{h_n}\right)\left(\frac{r}{h_n}\right)^l\right)^2 d F_{R \mid R \leq 0, D=0}(r) .
\end{aligned}
$$

In view of condition 1a holding, condition 3 holds if we have:

$$
\begin{aligned}
& \sup_{\rho(s,t)<\delta_n} \int_y (1(y \leq s)-1(y \leq t))^2 d F_{Y \mid R=r, R \leq 0, D=0}(y) \\
= & \sup_{\rho(s,t)<\delta_n} F_{Y \mid R=r, R \leq 0, D=0}(s)-2 F_{Y \mid R=r, R \leq 0, D=0}(s \wedge t)+F_{Y \mid R=r, R \leq 0, D=0}(t) \\
= & \sup_{\rho(s,t)<\delta_n} F_{Y \mid R=r, R \leq 0, D=0}(s \vee t)-F_{Y \mid R=r, R \leq 0, D=0}(s \wedge t) \rightarrow 0,
\end{aligned}
$$

for every $\delta_n \downarrow 0$. This holds under equicontinuity of $F_{Y|R=r,R\leq 0,D=0}$, which follows from the uniform boundedness of $f_{Y|R=r,R\leq 0,D=0}$ in A3.

Finally, by example 2.11.24 on page 221 of van der Vaart and Wellner, condition 4 is satisfied since $\mathcal{F}_n$ is VC class with a VC index of 2. To see this, note that every one-point set is shattered, but a two-point set:

$$
\{x_1, x_2\} = \left\{\begin{pmatrix} y_1 \\ r_1 \end{pmatrix}, \begin{pmatrix} y_2 \\ r_2 \end{pmatrix}\right\},
$$

with, say, $y_1 < y_2$ is not shattered because the function cannot pick out $\{x_2\}$. This establishes that the $B_{n,l}(y)$ terms converge. A similar argument applies to the $A_{n,l}$ terms. By the Cramér-Wold device the terms converge jointly. Finally, we need to establish the Hadamard differentiability of (1.19). Define (1.19) as a map $\phi: \left(\mathbb{R}^3 \times \ell^\infty(\mathbb{R})^2\right) \rightarrow \mathbb{R}$, and define $\ell^\infty(\mathbb{R})$ as the set of all uniformly bounded real functions on the real line, and $C(\mathbb{R})$ as the space of continuous functions on the real line. As a map from $\mathbb{R}^5$ to $\mathbb{R}$ (for a fixed value of y) the usual differentiability of ϕ implies Hadamard differentiability tangentially to the subspace $\mathbb{D}_0 = \left(\mathbb{R}^3 \times C(\mathbb{R})^2\right)$, and the conclusion follows by the functional delta method. ■

Proof of Theorem 3

Proof. The estimator (1.5) is a (Hadamard) differentiable function of several intermediate estimators. Let the vector of component quantities in (1.8) and (1.9) be

$$P=\begin{pmatrix} \lim_{r\to 0^+} F_{Y|D=1,R=r}(y) \\ \lim_{r\to 0^-} F_{Y|D=1,R=r}(y) \\ \lim_{r\to 0^-} F_{Y|D=0,R=r}(y) \\ \lim_{r\to 0^+} F_{Y|D=0,R=r}(y) \\ \lim_{r\to 0^+} E[D|R=r] \\ \lim_{r\to 0^-} E[D|R=r] \end{pmatrix},$$

with $\hat{P}$ as the corresponding vector of estimators. The j-th element of $\hat{P}$ is a local linear estimator of the conditional expectation of some variable, L_j , approaching $R=0$ from the right or the left. Hahn, Todd, and van der Klaauw (1999) established the joint convergence in distribution of local linear estimators of this type. Given assumptions A1-A7, which are analogous to their conditions, we can apply their result to establish the joint convergence in distribution of $\hat{P}$:

$$n^{\frac{1-b}{2}}\left[\hat{P}-P\right]\xrightarrow{d} N\left[\mathbf{0},\Sigma_P\right],$$

where Σ_P is a diagonal matrix with diagonal elements given by $\omega^+\sigma_1^{2+}(y)$, $\omega^-\sigma_1^{2-}(y)$, $\omega^-\sigma_0^{2-}(y)$, $\omega^+\sigma_0^{2+}(y)$, $\omega^+p^+(0)(1-p^+(0))$, and $\omega^-p^-(0)(1-p^-(0))$, respectively, and ω^+ and ω^- are defined by (1.13). The off-diagonal elements are zero since the first four estimators in $\hat{P}$ use separate observations, and they all involve conditioning either on $D=1$ or $D=0$. Choosing the bandwidth according to Assumption A7 in this section undersmooths, in the sense of Horowitz (2001b), causing the bias squared to converge to zero at a faster rate than the variance, correctly centering the asymptotic distribution.

Next I turn to the joint limiting distribution of the local linear conditional distribution function estimators, $\begin{pmatrix} \hat{F}_{Y_1|C,R=0}(y) \\ \hat{F}_{Y_0|C,R=0}(y) \end{pmatrix}$. This vector of estimators is a differentiable function of $\hat{P}$, so by the multivariate delta method and Lemma 2 we have that $\begin{pmatrix} \hat{F}_{Y_1|C,R=0}(y) \\ \hat{F}_{Y_0|C,R=0}(y) \end{pmatrix}$ converges as a process:

$$n^{\frac{1-b}{2}} \left[\begin{pmatrix} \hat{F}_{Y_1|C,R=0}(y) \\ \hat{F}_{Y_0|C,R=0}(y) \end{pmatrix} - \begin{pmatrix} F_{Y_1|C,R=0}(y) \\ F_{Y_0|C,R=0}(y) \end{pmatrix} \right] \xrightarrow{d} N\left[\mathbf{0}, J_P \Sigma_P J_P'\right],$$

where J_P is the Jacobian of the map from $\hat{P}$ to $\begin{pmatrix} \hat{F}_{Y_1|C,R=0}(y) \\ \hat{F}_{Y_0|C,R=0}(y) \end{pmatrix}$ in (1.8,1.9) evaluated at the truth, given by (1.16) in the text.

Finally I apply the functional delta method to derive the the limiting distribution of the quantile treatment effect estimator, (1.5). In terms of vectors, (1.5) can be written:

$$\hat{\delta}_{LQTE}(\tau) = c' \begin{bmatrix} \hat{F}^{-1}_{Y_1|C,R=0}(\tau) \\ \hat{F}^{-1}_{Y_0|C,R=0}(\tau) \end{bmatrix}, \tag{1.21}$$

which leads to the conclusion:

$$n^{\frac{1-b}{2}} \left(\hat{\delta}_{LQTE}(\tau) - \delta_{LQTE}(\tau)\right) \xrightarrow{d} N\left[0, c' J_{F_C} J_P \Sigma_P J_P' J_{F_C}' c\right],$$

where J_{F_C} is the Jacobian of the inverse in (1.21) evaluated at the true quantile, given by (1.15) in the text. This completes the derivation of the result. ■

The preceding theorem was for the simplest local linear distribution without smoothing in the y-direction. If we add smoothing in the y direction, then using Yu and Jones's (1998) Lemmas 1 and 2, a typical diagonal element of the variance covariance matrix of $\hat{P}$ is:

$$Var\left(\hat{P}_j|R\right) = \frac{R(K)}{nh_1 f_R(0)} \left(F(y|0)(1 - F(y|0)) - f_y(y|0)\,\alpha(W)\,h_2\right) + o_p\left(\frac{h_2^2}{nh_1}\right),$$

where $R(K) = \int K(u)^2 du$ and $\alpha(W) = \int \Omega(t)(1 - \Omega(t))\,dt$. If we continue to set the x-direction bandwidth as $h_1 = \gamma_{h_1} n^{-b}$ and $h_2 = \gamma_{h_2} h_1^2$, as suggested by Yu and Jones (1998) then the limiting variance becomes:

$$n^{1-b} Var\left(\hat{P}_j|R\right) = \frac{R(K)}{\gamma_{h_1} f_R(0)} F(y|0)(1 - F(y|0)) + O_p\left(\frac{1}{n^{2b}}\right),$$

and the $f_y(y|0)\,\alpha(W)\,h_2$ term drops out in the limiting variance, and the limiting variance derived in the proof continues to hold. The bias-squared term in the double smoothed estimator is:

$$B^2 = \left\{\frac{1}{2}F^{20}(y|0)\,\mu_2(K)\,h_1^2 + \frac{1}{2}F^{02}(y|0)\,\mu_2(W)\,h_2^2\right\}^2,$$

where $F^{ab}(y|z) = \frac{\partial^2 F(y|z)}{\partial z^a \partial y^b}$. Plugging in the rule for bandwidth, we get:

$$n^{1-b}B^2 = O_p\left(\frac{1}{n^{\frac{5b-1}{2}}}\right).$$

Since $b > 1/5$, the limiting bias is zero in the double-smoothed case as well, and thus under our conditions the limiting distribution for the double-smoothed estimator is the same as for the simpler estimator.

Appendix B: Bandwidth Selection

The estimation procedure I propose for estimating the p-th quantile treatment effect involves choosing two bandwidths for each conditional distribution function being estimated: $h_{1,p}$ for local linear smoothing along the running variable, R; and $h_{2,p}$ for smoothing "in the y-direction." Yu and Jones (1998) derive operational rules of thumb for choosing these bandwidths in the context of local linear quantile estimation:

$$\begin{aligned}
h_{1,p} &= h_{\text{mean}}\left\{p(1-p)/\phi\left(\Phi^{-1}(p)\right)^2\right\}^{1/5} \\
h_{2,p} &= \max\left(\frac{h_{1,1/2}^5}{h_1^3}, \frac{h_{1,p}}{10}\right) \text{ if } h_{1,1/2} < 1 \\
&\quad \text{and } \frac{h_{1,1/2}^4}{h_{1,p}^3} \text{ otherwise,}
\end{aligned}$$

where p is the quantile index being estimated, ϕ and Φ are the standard normal density and cdf, respectively, and h_{mean} is a suitable bandwidth for estimating a conditional mean at a boundary. I now turn to the choice of h_{mean}. I'll describe the case where we are estimating $\lim_{r\to 0^+} F_{Y|D=1,R=r}(y)$. The cases for estimating the other conditional distributions are analogous. Let $m(0) \equiv \lim_{r\to 0^+} E[Y|D=1, R=r]$ be the conditional mean corresponding to the conditional distribution we are estimating. As Imbens and Lemieux (2008) suggest for the RD context, I choose h_{mean} to minimize the approximate mean squared error of an estimator of this quantity. However, instead of the cross-validation method they propose, I suggest the plug-in method. This has been shown to have better convergence properties theoretically (see Ruppert, et al, 1995 for discussion), and in practice I have found plug-in to be more reliable than cross-validation. Fan (1992) shows that the mean squared error of this estimator can be written:

$$\begin{aligned}
& E\left[\left.(\hat{m}(0)-m(0))^2\right|\{R_i\}_{i=1}^n\right] \\
= \; & \frac{1}{4}\left(\alpha_K(0)\,m''(0+)\right)^2 h_n^4 + \frac{\beta_K(0)}{nh_n}\frac{\sigma_1^2(0+)}{f_{R|D=1,R\geq 0}(0)} + o_p\left(h_n^4+\frac{1}{nh_n}\right),
\end{aligned}$$

where we have $s_{l,c}=\int_{-\infty}^{c}K(u)\,u^l du \quad (l=0,1,2,3)$,

$$\begin{aligned}
\alpha_K(c) &= \frac{s_{2,c}^2 - s_{1,c}s_{3,c}}{s_{2,c}s_{0,c}-s_{1,c}^2}, \\
\beta_K(c) &= \frac{\int_{-\infty}^{c}\left(s_{2,c}-us_{1,c}\right)^2 K^2(u)\,du}{\left[s_{2,c}s_{0,c}-s_{1,c}^2\right]^2}.
\end{aligned}$$

Simple algebra shows that the optimal bandwidth is

$$h_n^{opt} = n^{-1/5}\left(\frac{\beta_K(0)\frac{\sigma_1^2(0+)}{f_{R|D=1,R\geq 0}(0)}}{\alpha_K^2(0)\left(m''(0+)\right)^2}\right)^{1/5}.$$

The plugin method (Hall, Sheather, Jones, and Marron, 1991) consists of estimating the quantities in this expression to compute the optimal bandwidth. The quantities to be estimated are $m''(0+)\equiv\lim_{r\to 0+}\frac{\partial^2}{\partial r}E[Y|D=1,R=r]$, $\sigma_1^2(0+)\equiv\lim_{r\to 0+}Var(Y|D=1,R=r)$. For $\sigma_1^2(0+)$ and $m''(0+)$ I adapt the methods of Ruppert, Sheather, and Wand (1995). The estimator for $m''(0+)$ is obtained by estimating a quartic fit in the $(R\geq 0, D=1)$ cell:

$$\hat{m}''(0+) = 2\mathbf{e}_3'\left(\mathbf{R}_{4,0}'\mathbf{R}_{4,0}\right)^{-1}\mathbf{R}_{4,0}'\mathbf{Y},$$

where $\mathbf{e}_s$ denotes a column vector with a one in the s-th position and zeros elsewhere,

$$\mathbf{R}_{p,c} = \begin{bmatrix} 1 & R_1-c & \cdots & (R_1-c)^p \\ \vdots & \vdots & \vdots & \vdots \\ 1 & R_n-c & \cdots & (R_n-c)^p \end{bmatrix},$$

and n is the number of observations in the cell. The estimator for $\sigma_1^2(0+)$ is obtained using the residuals from this same quartic fit:

$$\hat{\sigma}_1^2(0+) = (n-5)\sum_{i=1}^{n}\left\{Y_i - \hat{m}^Q(R_i)\right\}^2,$$

where $\hat{m}^Q(c)$ is a least squares quartic fit. Finally, for $f_{R|D=1,R\geq 0}(0)$ I use the following boundary kernel estimator (Jones, 1993):

$$\hat{f}_{R|D=1,R\geq 0}(0) = (nh)^{-1}\sum_{i=1}^{n} K_L\left(\frac{R_i}{h}\right),$$

where n is the sample size conditional on $D = 1$ and $R \geq 0$, and the boundary kernel is given by:

$$K_L(r) = \frac{(s_{2,0} - s_{1,0}r)\,K(r)}{s_{0,0}s_{2,0} - s_{1,0}^2},$$

and the bandwidth is chosen by a simple rule of thumb (Silverman's rule).

Plugging these estimates in, I compute the bandwidth for the mean as

$$h_{\text{mean}} = n^{-1/5}\left(\frac{\beta_K(0)\,\frac{\hat{\sigma}_1^2(0+)}{\hat{f}_{R|D=1}(0+)}}{\alpha_K^2(0)\left(\hat{m}''(0+)\right)^2}\right)^{1/5}.$$

In the case of the uniform kernel, $K(u) = .5 \times 1(|u| \leq 1)$, the constants in this expression are $\alpha_K(0) = 1/6$ and $\beta_K(0) = 4$.

Chapter 2

Union Wage-setting and the Distribution of Employees' Earnings: Evidence from Certification Elections

2.1 Introduction

How do unions affect the earnings distribution? This question is at the heart of the debate over the causes of increasing U.S. inequality over the past three decades. While market forces such as international trade and the supply of and demand for skilled labor have probably played a role (Juhn, Murphy, and Pierce, 1993; Katz and Murphy, 1992), institutional forces such as falling unionization rates may also have contributed (Freeman, 1993; Blau and Kahn, 1996; DiNardo, Fortin, and Lemieux, 1996; Card, 2001). From 1979 to 2009 the U.S. private sector unionization rate fell from about 25 percent to 8 percent (Bureau of Labor Statistics, 2010; DiNardo, Fortin, and Lemieux, 1996). To estimate the impact of this drop in the unionization rate, DiNardo, Fortin, and Lemieux (1996) construct counterfactual wage densities based on observed characteristics and show unionization is associated with substantial wage compression. If this compression reflects the causal effect of unionization, then deunionization accounts for a significant part of the increase in U.S. earnings inequality.

Comparisons between the earnings of unionized and non-unionized workers robustly show a positive union wage gap, especially in lower skill groups, but recent efforts to estimate the causal effect of unionization have generated mixed results.[1] Quasi-experimental evidence from DiNardo and Lee (2004), for example, shows little effect on employer outcomes, apparently at odds with regression-based comparisons. At the same time, this small average effect may mask significant, but offsetting effects on different features of the distribution.

The main objective of this paper is to identify the causal effect of unionization on the distribution of employee earnings. The target of estimation can be understood in terms of a hypothetical experiment where a set of establishments are randomly assigned to be unionized or not. The causal effect of interest is the difference between the subsequent distribution of earnings among employees at the unionized and non-unionized establishments. To approximate this hypothetical experiment, this paper adapts DiNardo and Lee's (2004) regression discontinuity (RD) design based on union certification elections, using administrative records on individual earnings matched to establishment-level election results. If establishments where the union barely won and barely lost are otherwise comparable, the resulting difference in the distribution of employees' earnings is due to the causal effect of unionization.

The paper motivates the possibility of distributional effects by analyzing the union wage-setting problem in the context of a certification election. The theoretical section of the paper highlights how the the need to garner the support of a majority of workers can lead to a wage schedule that raises the wages of lower-productivity workers but reduces the wages of higher-productivity workers, resulting in a wage-compressing effect of unionization.

Consistent with DiNardo and Lee (2004), the RD estimates reported here show little effect of unionization on average earnings. At the same time, my results provide clear evidence of a distributional effect. Specifically, unionization raises the lower tail of the earnings distribution by around 25 log points, while reducing earnings at the very high end. Further empirical results show unionization increases turnover for higher productivity workers, but not for lower productivity workers, consistent with the model's interpretation of union wage compression as reflecting a wage schedule that raises lower-productivity workers' wages but reduces the return to skill.

[1] Lewis (1986) surveys the large early literature on union wage gaps, while Blanchflower and Bryson (2003) offer an updated look using newer data. Quasi-experimental evidence on the causal effect of unionization includes DiNardo and Lee (2004) and Lee and Mas (2009).

The remainder of the paper proceeds as follows. The next section describes the U.S. private sector unionization process and highlights the theoretical implications this process has for the distribution of earnings. Section 2.3 describes the data used in the empirical work. Section 2.4 lays out the research design and the econometric framework for identifying and estimating the effect of unionization on the distribution of earnings, and section 3.4 presents the estimation results. Section 2.6 summarizes the findings and concludes.

2.2 Background

2.2.1 NLRB election process

Since 1935, most U.S. private sector unionization has been governed by the National Labor Relations Act (NLRA), which specifies the rights of unionizing workers. While an employer may voluntarily bargain with the workers' chosen representative, or in some cases may be required to do so even without an election, the traditional process by which workers unionize is through a National Labor Relations Board (NLRB) secret ballot election.[2] Although in practice an organizing drive is often fraught with disputes and delays, the following steps describe the stylized path a group of workers follows to form a union:[3]

1. Petition drive: Union organizers lobby workers, collect signatures expressing a desire to hold an election, and submit a petition to the NLRB to hold an election. If the petition is accepted, the NLRB ascertains the scope of the bargaining unit and sets the election time and place, usually the workplace.

2. Election: Eligible workers vote for or against the union, and the union wins if it receives a simple majority (50 percent + 1) of the votes cast.

3. Certification: If the union wins, the NLRB certifies it as the sole authorized representative of the workers in the bargaining unit, and requires the employer to bargain "in good faith" with the union.

[2]Secret ballot election has historically been the dominant form of new unionization, although in recent years voluntary recognition through neutrality agreements and card checks have become more common. (Brudney, 2005)

[3]The simple process laid out here follows the procedures described in NLRB (2010). See Ferguson (2008) and DiNardo and Lee (2004) for a more complete description of the possible complications and objections that can be raised at each step.

4. Bargaining: The employer negotiates with union representatives over a collective bargaining agreement. If an agreement is reached, the contract becomes binding for all employees in the unit.

NLRB certification elections may include two or more competing unions on the ballot. In the case of multiple competing unions, a simple majority is still required for certification. Elections may also be held to remove union representation altogether (decertification) or to replace one union with another. These cases, however, occur relatively infrequently, and the analysis focuses on certification elections.

2.2.2 Theoretical framework

How do the incentives a union faces in a certification election affect wage setting? The model developed here takes as its starting place that the primary objectives of union leaders are survival and expansion of the organization, and retention of their offices (Atherton, 1973; Ross, 1948; Berkowitz, 1954; Ashenfelter and Johnson, 1969). Accordingly, a union facing a certification election will pursue a wage agreement that maximizes the probability of winning. The theoretical framework in this section is similar to Farber (1978), Booth (1995), Acemoglu, Aghion, and Violante (2002), and Lee and Mas (2008), who also consider the effects of majority-rules politics on union wage policies, but I focus on the implications for the distribution of workers' wages.

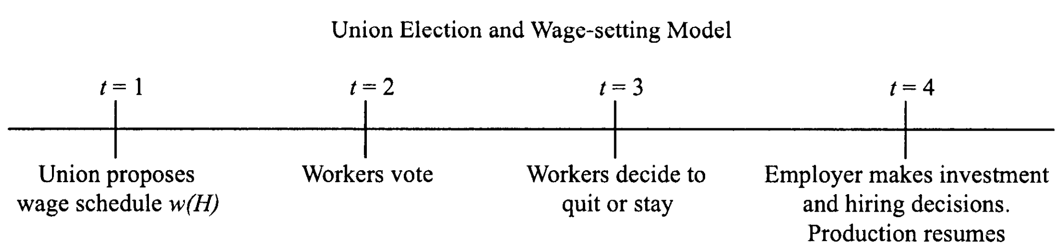

Figure 2-1: The model

The model casts the union election and wage-setting process as a four-stage game. The stages, shown in Figure 2-1, follow the stylized union certification process described above. In the first stage, a union petitions to represent the workers at a plant (currently producing in the competitive sector) and proposes a wage schedule, $w(H) = (v + r(H))H$, which gives a worker's wages as a function of his or her human capital, H. The outside (competitive) price of human capital is v, and $r(H)$ denotes the union rent earned by a worker with human capital H. The term "union rent" hightlights that $r(H)$ is the difference between the competitive and union price of skill. In the second stage, workers vote for the union if their union rent, $r(H)$, exceeds their individual cost of union representation, η, provided the wage schedule doesn't cause the plant to shut down. The voting rule can therefore be written $1(\eta \leq r(H))$, subject to the plant not shutting down. The cost of union representation η reflects any pecuniary (e.g., union dues) or nonpecuniary factors affecting workers' preferences for union representation outside of wage differences, and is assumed to be independent of H. In the third stage, workers decide whether to quit or stay, after observing the outcome of the election. Workers stay if their union rent exceeds their cost of unionization, net of an individual-specific switching cost, ε, also independent of H. The worker's decision rule after a union victory is therefore $1(\eta - \varepsilon \leq r(H))$, where the worker stays if the indicator is equal to one. Finally, the employer makes investment and hiring decisions to maximize profits.

The production technology combines H with another factor, K, in fixed proportions with constant returns to scale, so the production function can be written $Y(H) = yH$. Normalizing the H/K ratio to unity, the competitive return to K is $y - v$. A fraction ϕ_K of K is sunk in the production relationship.

Two assumptions in the model setup are important for the results. The first is that workers' decisions to vote for a union or quit their job depend on more than simply a comparison of wages. These other factors are modeled as union costs η and switching costs ε. The presence of the union cost in the model prevents the union from being able to assure 100 percent vote share with only an infinitesimal wage increase to all workers. The presence of the switching cost prevents the unrealistic scenario that all workers who vote against the union quit after a union victory, and thus widens the scope for the union to cater to specific groups of workers. The union and switching costs are assumed to reduce the effective union and outside return to human capital, respectively. While the specific functional form is a simplification, the substantive assumption that the "equivalence premium"—the dollar amount by which a union would have to raise a worker's earnings to make him indifferent between a union and no union—is on average increasing in a worker's outside wage is important. A concrete motivating example is that union dues are commonly collected as a percentage of wages. The consequence of this assumption for the model is that it makes it more efficient for the union to shift resources to attract the votes of workers with lower outside options.

The second key assumption is that employers face short term rigidities in adjusting inputs. The model captures this in a simple way following Caballero and Hammour (1998) by assuming the chosen technology takes fixed ratios of inputs. The assumption is that while technology can adjust over time, it is essentially fixed over the period relevant for an initial collective bargaining agreement. Consequently, employers cannot undo the effects of a union wage schedule by immediately adjusting the production inputs. An alternative way of capturing this that leaves the results substantively unchanged is to allow the employer to fully adjust K, but have the union impose firing costs.

The union's optimal wage schedule is shown in the theory appendix to maximize the expected vote share, subject to a minimum profit constraint for the employer:

$$\max_{r(h)} \int F_\eta \left(r\left(h\right)\right) dF_H \left(h\right)$$
$$\text{s.t.} \quad \int \left(\phi_K \left(y-v\right) - r\left(h\right)\right) h F_{\eta-\varepsilon} \left(r\left(h\right)\right) dF_H\left(h\right) \geq 0.$$

Concrete functional forms illustrate the main implications of the model. Assume H is log-normally distributed and η and ε are exponentially distributed with parameters λ_η and λ_ε, where the mean switching cost, $1/\lambda_\varepsilon$, exceeds the average rent, $\phi_K\left(y-v\right)$. In this case the optimal rent schedule is

$$r\left(H\right) = \begin{cases} r^+\left(H,\lambda^*\right) & , \quad H \leq h_1\left(\lambda^*\right) \\ -\left(\frac{1}{\lambda_\varepsilon} - \phi_K\left(y-v\right)\right) & , \quad H > h_1\left(\lambda^*\right) \end{cases}, \tag{2.1a}$$

where $r^+\left(H,\lambda\right)$ satisfies

$$\exp\left(\lambda_\eta r\right) = \frac{\lambda_\eta}{\lambda H} + \frac{\lambda_\varepsilon}{\lambda_\eta + \lambda_\varepsilon}\left(1 + \left(\phi_K\left(y-v\right) - r\right)\lambda_\eta\right). \tag{2.1b}$$

The threshold above which rents are negative is

$$h_1\left(\lambda\right) = \frac{1}{\lambda}\frac{\lambda_\eta + \lambda_\varepsilon}{1 - \lambda_\varepsilon \phi_K\left(y-v\right)}, \tag{2.1c}$$

and λ^* is the value of λ that satisfies the profit constraint with equality.

The optimal union wage schedule, (2.1), has two features with stark implications for the union effect on the distribution of earnings. First, the union rent is positive for lower-productivity workers, and negative for higher-productivity workers. Interestingly, the less the union can extract from the employer (i.e., the smaller is ϕ_K), the lower the threshold above which rents are negative. The second feature is that even where rents are positive, they are decreasing in the level of human capital:

$$\left.\frac{dr}{dH}\right|_{H=h\leq h_1} = -\left(\lambda h^2\left(\frac{\lambda_\varepsilon}{\lambda_\eta+\lambda_\varepsilon}+\exp\left(\lambda_\eta r\right)\right)\right)^{-1} < 0.$$

Thus compared to the competitive equilibrium, the union wage schedule compresses the distribution of potential wages and shifts it to the right.

To see the effects on the distribution of earnings directly, I solve a numerical example based on the distribution of wages observed in the the sample of full-time, nonunion workers from the 1998-2000 Current Population Survey (CPS). Taking the hourly wage as a measure of human capital, H, I assume the sunk fraction of K is $\phi_K = .3$. Normalizing the competitive return to human capital to be $v = 1$, and assuming labor's share in income is about .7, I set the production function parameter to be $y = 10/7$. Finally, I assume unionization and switching costs both have means of $1.5 \times \phi_K(y - v)$.

The optimal wage schedule in this example distributes rents disproportionately to low-wage workers and substantially compresses the distribution of earnings. Figure 2-2 plots the union premium (in dollars), $r(H) \times H$, implied by the optimal union wage schedule as a function of human capital. The union premium is positive for lower levels of human capital, and decreasing and eventually negative for higher levels of human capital. Figure 2-3 compares the distribution of potential union and non-union log wages implied by this example. The union distribution is compressed and shifted to the right relative to the non-union distribution.

Given the assumption of fixed K, these implications are likely to apply to the short term. In the long term, however, K can adjust, and the employer may be able to terminate lower-productivity workers, who are being paid above marginal product. In the longer term, therefore, the degree to which the union wage schedule inflates the wages at the bottom of the distribution relative to the top will be attenuated.

The theoretical discussion shows that unions facing certification elections have an incen-

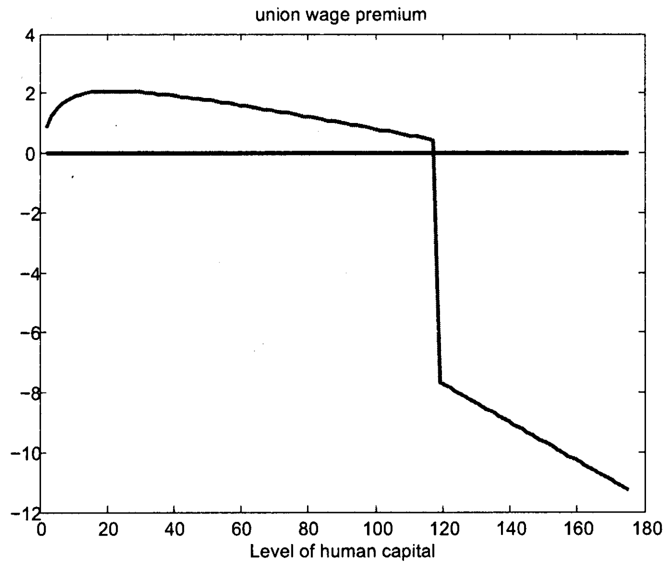

Figure 2-2: Union wage premium as a function of human capital for parameter values $\phi_K = .3$, $v = 1$, $y = 10/7$, $\lambda_\eta = \lambda_\varepsilon = (1.5\phi_K (y - v))^{-1}$.

tive to commit to a wage schedule that favors lower-productivity workers at the expense of higher-productivity workers. In the example, the union wage schedule raises the wages of low-productivity workers and lowers the wages of high-productivity workers, compressing the distribution of wages.

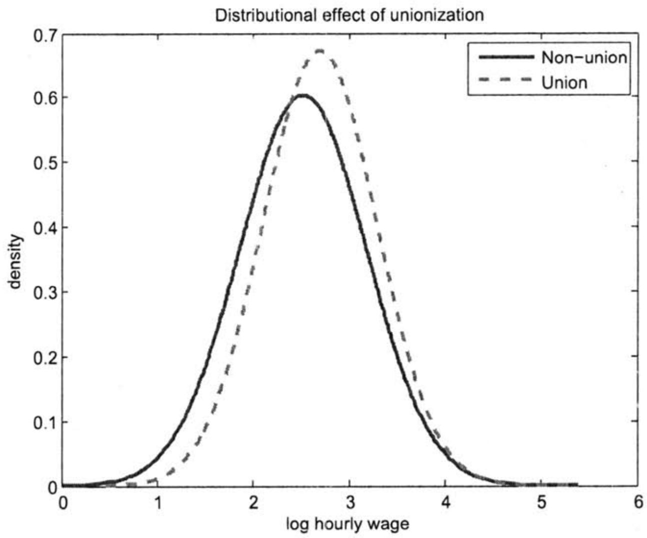

Figure 2-3: Densities of potential union and non-union log wages implied by the optimal wage schedule from the example in the text.

2.3 Data

2.3.1 Union Elections

The analysis uses a dataset on the universe of NLRB union representation election results from 1963 to 2006, which was compiled and analyzed by Ferguson (2008). Each record in this dataset represents a union certification election held at an establishment, and includes the number of votes cast for and against union representation, the date of the election, and the employer's name and address. The data appendix explains in detail how the dataset was constructed.

The main sample used in the analysis covers the years 1992-2001, the period covered by the earnings data described below. This sample contains data on 37,354 representation elections, involving over 1.7 million votes cast. Unions received 50.4 percent of the votes cast, and won 54.5 percent of the elections.[4] Figure 2-4 shows the distribution of the union vote share in the sample. The mode is around 40 percent, with a significant number of elections in which the union received all votes. The figure corresponds to Figure II in DiNardo and Lee (2004), who use similar data. Figure 2-5 shows the distribution of the union margin of victory in terms of number of votes, close to the threshold. This figure shows that close

[4]The votes-weighted union success rate is only 41.5 percent, indicating that unions fare considerably worse at large establishments, as noted by Farber (1999).

elections represent the typical case, a fact that is important for the interpretation of the estimation results.

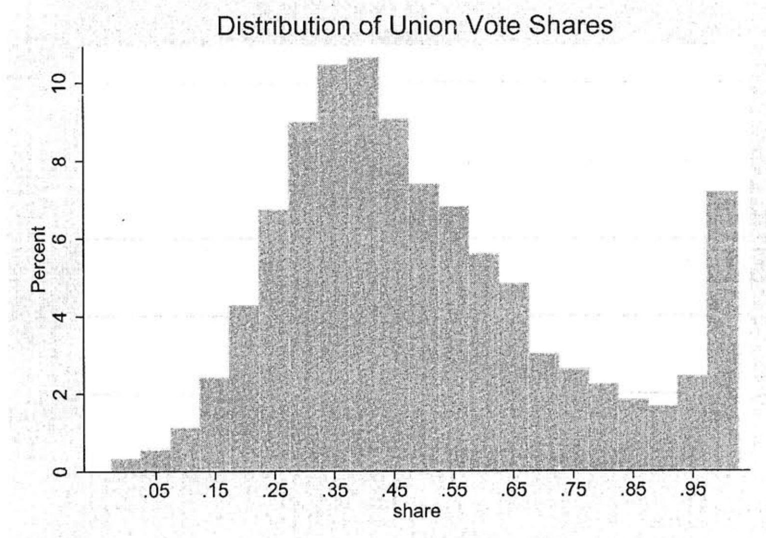

Figure 2-4: Votes-weighted histogram of the union vote share in representation elections from 1992 to 2001. Data are from NLRB election records, restricted to elections where 10 or more votes were cast.

2.3.2 LEHD

The second data component contains individual-level earnings from the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) database. The LEHD combines data from a wide variety of state and federal administrative records and surveys. In particular, the LEHD integrates the universe of unemployment insurance-covered (UI) earnings records held by participating state agencies into a cohesive data structure using person and employer identifiers, allowing linkages to other sources of data.[5]

The Employment History Files (EHF) within the LEHD contain quarterly records of individuals' UI-covered earnings. The EHF for each of the 23 covered states contains a record for each employee-employer combination—a job—that produced at least one dollar of wages in that state in each year. The data cover a period as wide as 1985 to 2004, although for most states the data only go back to the early 1990s. The EHF contains more

[5]For more details on the construction and uses of the LEHD database, see McKinney and Vilhuber (2008), Lane (2008), Abowd, Haltiwanger, and Lane (2004), and Abowd, Stephens, Vilhuber, Andersson, McKinney, Roemer, and Woodcock (2009).

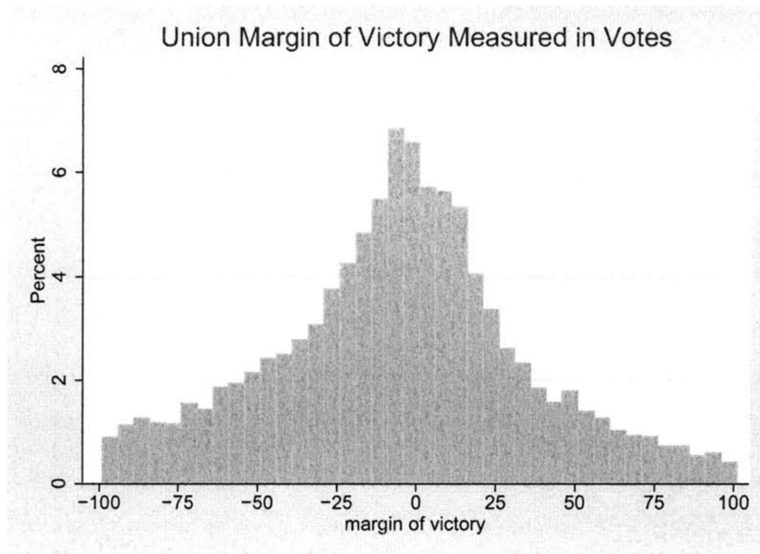

Figure 2-5: Votes-weighted histogram of the union margin of victory (in terms of number of votes) in representation elections from 1992 to 2001. Only elections decided by 100 or fewer votes are included in the histogram. Data are from NLRB election records, restricted to elections where 10 or more votes were cast.

than 2.8 billion records, although I focus on full-time, full-year workers with high labor force attachment. This is the sample for whom earnings most closely approximates the hourly wage, on which the theory is based. Restricting to this sample also allows us to separate wage effects from labor supply effects. See the data appendix for details on the sample selection.

Crucially for this study, individual-level earnings records in the LEHD can be matched to establishments. For each union election record in the NLRB election dataset, employees at the time of the election can be identified by matching employer name and address information from the election record with employer information in the LEHD. The data appendix describes in detail the procedure used for matching the two datasets. The matching procedure identified over 1.5 million individuals who were employed at establishments at the time a union election was held from 1992-2001. The subsequent earnings and employment histories of these individuals constitute the main outcomes of interest.

Table 2.1 reports summary statistics on pre-election earnings, post-election earnings, and retention by union status for full-time, full-year workers in my sample. Average (post-election) annual earnings—defined as the sum of the four quarterly earnings beginning six months after the election—in this sample is just under $30,000. Unionized workers earn on average nearly \$3,300 more than non-unionized workers, consistent with the large literature on union wage gaps finding a union premium of about 10 percent. This wage gap, however, reflects both the causal effect of unionization, as well as a selection effect. The nearly \$5,000 pre-election earnings gap suggests the selection effect is large. The table also reports statistics for the sample restricted to close elections (+/- 5 percent). Most of the earnings differences disappear in the restricted sample, consistent with bias in the full-sample comparisons by union status, and with DiNardo and Lee's (2004) findings of a small average effect of unionization. However, this small average effect could be masking significant, but offsetting effects elsewhere in the distribution. Average earnings also appear to be lower in the sample restricted to close elections, which reflects that the largest (and highest-paying) employers are less likely to be involved in close union elections.

Table 2.1: Pre-election Earnings and Outcomes by Unionization Status

	Full Sample			Discontinuity Sample		
	All	Non-unionized	Unionized	All	Non-unionized	Unionized
Annual earnings	29,299	27,319	30,614	22,559	22,306	22,871
	(41,966)	(46,565)	(38,556)	(42,546)	(50,493)	(29,911)
Pre-election earnings	43,202	40,129	45,065	34,474	34,884	34,028
	(47,082)	(46,173)	(47,528)	(38,862)	(47,543)	(26,351)
Stayed employed	0.523	0.505	0.535	0.48	0.492	0.465
at plant	(.499)	(.500)	(.499)	(.500)	(.500)	(.499)
N	1,539,325	614,159	925,166	145,739	80,618	65,121

Notes: the table reports sample means and standard deviations of the dependent variables at left by union status for the sample of full-time, full-year workers employed at plants in covered states where a union election was held from 1992-2001. Earnings variables are measured in 2000 dollars. Annual earnings are defined as the sum of the four quarterly earnings starting two quarters after the union elections. Pre-election earnings are the sum of the four quarterly earnings prior to the union election closing date. Stayed employed at plant is as of 10 quarters after the union election closed. The discontinuity sample restricts to elections that were decided by 5 percent or less.

2.4 Research Design and Econometric Framework

A fundamental obstacle to measuring the effect of unionization on earnings is selection bias: earnings within unionized plants may differ for reasons other than union representation. This study seeks to overcome selection issues by using a regression discontinuity (RD) research design, originally developed by Thistlethwaite and Campbell (1960). The motivation for the design, first used in this context by DiNardo and Lee (2004), is that new unions arise through a majority-rule election. If plants and workers where the union barely won and barely lost are comparable, then close union elections approximate a randomized experiment, and the resulting difference in the distribution of earnings provides a reliable estimate of the causal effect of unionization.[6] To represent this idea formally, let $D = 1(R > 0)$ be an indicator for union representation, where R is the union margin of victory (negative for losses). Let Y_1 be an individual's earnings under union representation, and let Y_0 be the earnings otherwise, so that observed earnings is $Y = Y_0 + (Y_1 - Y_0)D$.

Since a worker is never observed simultaneously with and without union representation, we cannot measure the individual specific treatment effect, $Y_1 - Y_0$, but we can estimate the treatment effect on the distribution of outcomes, that is, the difference between the distributions of Y_1 and Y_0 at the margin of union victory. The distributional effects of unionization are captured by the quantile treatment effect, or the difference between the quantiles of potential earnings:

$$\delta(\tau) \equiv Q_{Y_1|R=0}(\tau) - Q_{Y_0|R=0}(\tau). \tag{2.2}$$

Note that this is the treatment on the distribution of earnings as a whole, rather than the effect of treatment on any particular individual. While (2.2) does not capture the effect of unions that win or lose by large margins, Figure 2-5 implies that the effect conditional on a close election reflects the typical case.

The key identifying assumption is that the conditional distribution of potential earnings as a function of the union vote share is smooth near the threshold of union victory, and thus any jumps in the observed distribution of earnings at the threshold is due to the treatment. Formally:

Assumption 1: Local Smoothness $F_{Y_d|R}(y|r)$ is continuous in r over an ε-neighborhood of zero, and is strictly increasing in y over the same neighborhood, for $d \in \{0, 1\}$.

This assumption is satisfied if, for example, unions, workers, and firms are *a priori* uncertain about the outcome of the election when it is close (see Lee, 2008 for a formal proof). The condition that the distribution be increasing in y ensures that quantiles are uniquely defined at the threshold. Given the sharp RD design setup and this local smoothness assumption, the identifying conditions given in Frandsen (2008) are satisfied.[7]

[6] See Lee (2008) for further discussion on the conditions under which a close election provides as-good-as-randomized variation.

[7] The other identification conditions in Frandsen (2008)—that the probability of treatment jumps discretely at the threshold, and that crossing the threshold has a montonic effect on treatment status—are automatically satisfied in the sharp RD setup considered here.

I focus here on the effect of union representation, which leads to the sharp RD design (Campbell, 1969). Another possible treatment of interest is an indicator for a collective bargaining agreement, leading to a "fuzzy" design since an agreement does not always follow from a union victory (Ferguson, 2008). Data on whether an agreement was reached and a contract signed can be inferred from Federal Mediation and Conciliation Service (FMCS) records on contract expiry and renewal. While the econometric framework described in the text applies equally well to the sharp and fuzzy designs, studying the effect of a collective bargaining agreement in this context suffers from two problems. The first is that FMCS records severely undercount union agreements (DiNardo and Lee, 2004), introducing a potential source of bias. Second, there is typically a time lag of several months between a certification election and a collective bargaining agreement being reached. Any responses by employees or employers to the outcome of the election, but prior to an agreement being reached, contaminates the design and would lead to a discontinuity in the distribution of Y_0 (potential earnings under no collective bargaining agreement) at the margin of union victory. For these reasons, the current research design is ill-suited for studying the effect of collective bargaining agreements, and I focus on the effect of union representation.

Frandsen (2008) outlines a tractable procedure for estimating quantile treatment effects in this framework. In the sharp RD design considered here, the estimator becomes particularly simple: it is the difference between kernel-smoothed local linear estimates of the quantiles of earnings approaching the threshold of union victory from the right and from the left. Formally, the estimator can be written:

$$\hat{\delta}_{LQTE}(\tau) = \hat{F}^{-1}_{Y_1|R=0}(\tau) - \hat{F}^{-1}_{Y_0|R=0}(\tau),$$

where

$$\begin{aligned} \hat{F}^{-1}_{Y_1|R=0}(\tau) &= \inf\left\{a : \hat{F}_{Y_1|R=0}(a) = \tau\right\}, \\ \hat{F}^{-1}_{Y_0|R=0}(\tau) &= \inf\left\{b : \hat{F}_{Y_0|R=0}(b) = \tau\right\}, \end{aligned}$$

and $\hat{F}_{Y_1|R=0}(y)$, $\hat{F}_{Y_0|R=0}(y)$ are local linear, consistent estimates of the conditional distribution functions of potential earnings at the threshold,

$$\hat{F}_{Y_1|R=0}(y) = \frac{1}{\sum_{j:D=1} w_j(h_1)} \sum_{j:D=1} w_j(h_1)\,\Omega\left(\frac{y-Y_j}{h_2}\right),$$
$$\hat{F}_{Y_0|R=0}(y) = \frac{1}{\sum_{j:D=0} w_j(h_1)} \sum_{j:D=0} w_j(h_1)\,\Omega\left(\frac{y-Y_j}{h_2}\right).$$

The weighting function associated with local linear fitting is given by:

$$w_j(h_1) = K\left(\frac{R_j}{h_1}\right)\left[S_{n,2} - R_j S_{n,1}\right],$$

with

$$S_{n,l} = \sum_{i:D=0,Z=0} K\left(\frac{R_i}{h_1}\right) R_i^l, \quad l = 1, 2.$$

The bandwidths h_1 and h_2 are chosen to minimize the approximate mean squared error; $K(\cdot)$ is a kernel density fucntion; and $\Omega(\cdot)$ is a kernel distribution function.

The estimator for the local quantile treatment effect is asymptotically normally distributed with the following limiting distribution:

$$n^{\frac{1-b}{2}}\left(\hat{\delta}_{LQTE}(\tau) - \delta_{LQTE}(\tau)\right) \xrightarrow{d} N\left[0, \frac{\omega^+\tau(1-\tau)}{f_{Y_1|R=0}\left(Q_{Y_1|R=0}(\tau)\right)} + \frac{\omega^-\tau(1-\tau)}{f_{Y_0|R=0}\left(Q_{Y_0|R=0}(\tau)\right)}\right],$$

where the bandwidths h_1 and h_2 are proportional to n^{-b}, with $b \in (1/5, 1)$, and $f_{Y_1|R=0}$ and $f_{Y_0|R=0}$ are the densities of Y above and below the threshold, respectively. The other constants are given by:

$$\omega^+ = \frac{\int_0^\infty \left(s_2^+ - s_1^+ u\right)^2 K(u)^2\, du}{f_R(0)\,\gamma_{h_1} \cdot \left[s_2^+ s_0^+ - \left(s_1^+\right)^2\right]^2},$$
$$\omega^+ = \frac{\int_{-\infty}^0 \left(s_2^- - s_1^- u\right)^2 K(u)^2\, du}{f_R(0)\,\gamma_{h_1} \cdot \left[s_2^- s_0^- - \left(s_1^-\right)^2\right]^2},$$

where $s_l^+ = \int_0^\infty K(u)\, u^l du$ and $s_l^- = \int_{-\infty}^0 K(u)\, u^l du$. In the empirical results, I estimate the distribution of the estimator via the nonparametric bootstrap. The validity of the bootstrap in this setting is discussed in Frandsen (2008).

A regression discontinuity estimator for the average treatment effect at the threshold is given by:

$$\hat{\delta}_{ATE} = \hat{E}\left[Y_1|R=0\right] - \hat{E}\left[Y_0|R=0\right], \tag{2.3}$$

where

$$\begin{aligned}
\hat{E}\left[Y_1|R=0\right] &= \frac{1}{\sum_{j:D=1} w_j\left(h_1\right)} \sum_{j:D=1} w_j\left(h_1\right) Y_j, \\
\hat{E}\left[Y_0|R=0\right] &= \frac{1}{\sum_{j:D=0} w_j\left(h_1\right)} \sum_{j:D=0} w_j\left(h_1\right) Y_j,
\end{aligned}$$

and the weighting functions and bandwidths are as described above.

2.5 Results

Comparisons of earnings by union status suggest that unionized workers' earnings are higher on average and more compressed than non-unionized workers' earnings. These findings can be seen in Table 2.2, panel A, which reports OLS and quantile regression coefficients from a regression of post-election log earnings on an indicator for union representation status for the sample of full-time, full-year workers at establishments where 10 or more votes were cast in a union election. The first column in panel A shows the estimated difference in log earnings is on average .146 with a standard error of .0024. The remaining columns in panel A report differences in the sample quantiles of log earnings. The tenth percentile of unionized earnings is .1931 (s.e.=.0068) log points higher than the tenth percentile of non-unionized earnings. The difference in median earnings is .1667 (s.e.=.0023), and the difference in the 90th percentile is .0765 (s.e.=.0017). These estimates are consistent with regression-based comparisons, such as those surveyed by Lewis (1986), which robustly find a significant positive union wage gap, and quantile regression estimates (e.g., Chamberlain, 1994), which show the union-nonunion wage differential is monotonically declining in wage percentile.

Are these differences due to the causal effect of unionization? Regression discontinuity estimates support the notion that unionization has little effect on the average, but significantly compresses the distribution of employee earnings. These findings can be seen in Panel B of Table 2.2, which reports RD estimates of the average treatment effect, (2.3), in the first column, and estimates of the quantile treatment effect, (2.2), in the remaining

Table 2.2: Union Log Earnings Effect

	Quantiles				
Average	0.1	0.25	0.5	0.75	0.9
	A. OLS / Quantile regression				
.1464	.1931	.1679	.1667	.1341	.0765
(.0024)	(.0068)	(.0032)	(.0023)	(.0016)	(.0017)
	B. Regression Discontinuity				
.0015	.2472	.0701	.0749	.0564	.0447
(.0124)	(.0855)	(.0267)	(.0113)	(.0234)	(.0243)

Notes: The table reports estimates of the union effect on post-election log earnings for full-time, full-year workers in covered states from 1992-2001. Panel A shows full sample OLS and quantile regression estimates, and Panel B shows regression discontinuity estimates. The bandwidth (in terms of union vote share) is .025. Earnings (in 2000 dollars) are defined as the sum of the four quarterly earnings starting two quarters after the union election closed. Only workers at plants where more than 10 votes were cast in a union election are included.

columns. The first column shows the average effect of unionization conditional on a close election is small: .0015 with a standard error of .0124. This result is consistent with DiNardo and Lee's (2004) finding of little union effect on average wages. This small average effect masks larger effects elsewhere in the distribution, as the remaining columns in Panel B show. The effect on the 10th percentile is large and positive, .2472 with a standard error of .0855, while the effect on the median is a more modest .0564 (s.e.=.0234), and the effect on the 90th percentile is smaller still at .0447 (s.e.=.0243). The distributional effects of unionization are summarized graphically in Figure 2-6, which plots estimates and pointwise confidence intervals for quantile treatment effects for quantile indices from .1 to .9. The figure shows unionization significantly raised the lower end of the distribution, but had more modest effects through the middle and upper end of the distribution.

Motivated by the model's prediction that unions may reduce the highest productivity workers' wages, I take a closer look at the union effect on the upper tail of the earnings distribution. The results suggest that unionization significantly reduces the highest quantiles of employee earnings. These findings are summarized in Figure 2-7, which plots estimates and confidence intervals for the effect of unionization on the 90th through the 99th percentile of earnings. While the effect on the 90th percentile is small but positive, the confidence interval begins to include negative values around the 95th percentile, and the estimated

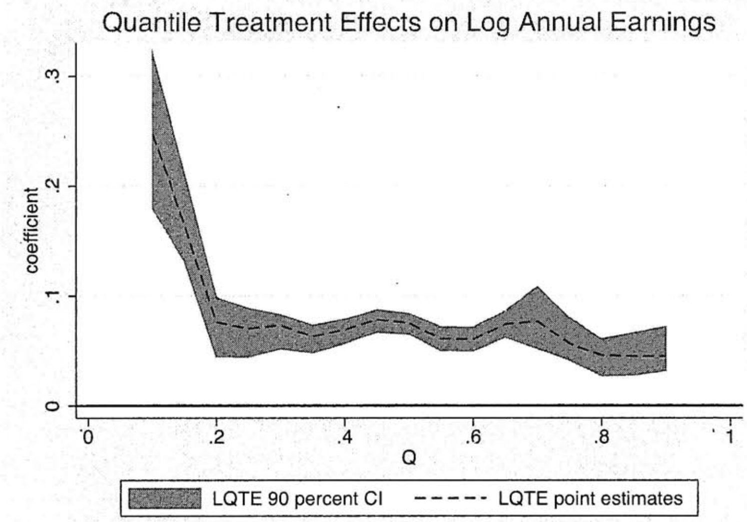

Figure 2-6: Estimates and 90-percent confidence intervals for the effect of unionization on the quantiles of employees' log annual earnings.

effect becomes large and significantly negative starting with the 97th percentile, reaching -.54 at the 99th percentile. In terms of magnitude, unionization's largest impact on earnings appears to be to cut off the upper tail.

The results thus far imply that unionization compresses and shifts the earnings distribution slightly to the right. This effect can be seen directly in Figure 2-8, which plots estimates of the counterfactual earnings densities by union status, conditional on a close election. The solid curve represents the density of potential earnings without unionization, and the dashed curve represents the density of potential earnings under unionization. The union density is lower in the tails, reflecting compression, and shifted slightly to the right, reflecting the modest positive effect throughout much of the distribution seen in Figure 2-6.

It is tempting to interpret the estimates plotted in Figures 2-6 and 2-7 as giving the effect of unionization on an individual of a given rank in the earnings distribution. In order to make this interpretation, one would have to invoke the rank invariance assumption that a worker with rank τ in the non-unionized potential earnings distribution also has rank τ in the unionized distribution (Heckman, Smith, and Clements, 1997; Chernozhukov and Hansen, 2005). In this case, the τ-th quantile treatment effect—the difference between the τ-quantile of the unionized and non-unionized potential outcome distributions—corresponds to the effect of unionization on a worker of rank τ, since that workers rank is unchanged

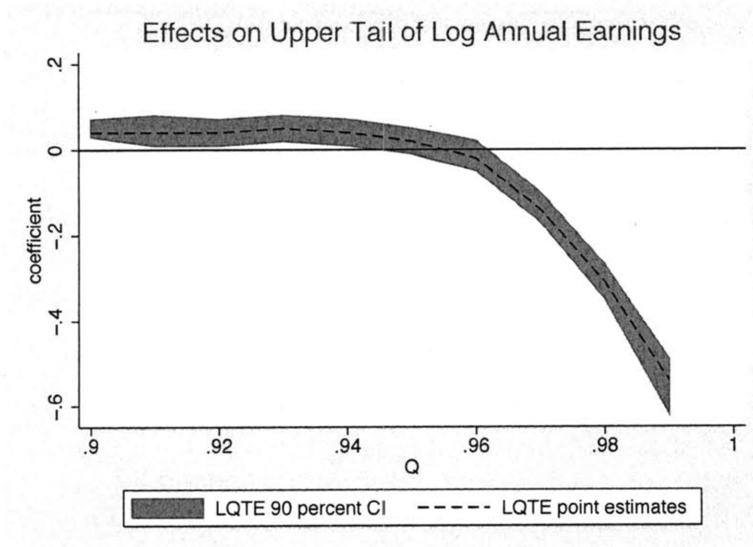

Figure 2-7: Estimates and 90-percent confidence intervals for the effect of unionization on the upper tail of employees' log annual earnings.

by unionization. Formally, this assumption could be written: $F_{Y_1|R=0}(Y_1) = F_{Y_0|R=0}(Y_0)$ almost surely. However, given the model's suggestion that unions may increase the pay of some workers at the expense of others, and the empirical results that union effects on quantiles vary drastically, this assumption may be unrealistic in this setting. For example, if unionization led to the termination of some higher-paid workers (e.g., management), ranks would almost certainly be affected.

An alternative assumption that may be more plausible, and allows for inference on the union effect on individuals across the distribution, is that an individual's rank in the non-unionized potential earnings distribution is equal to his or her rank in the pre-election earnings distribution. This would be true if, for example, a worker's rank remains unchanged from one year to the next, barring changes in union status. Denoting the distribution of pre-election earnings by $F_{Y_{-1}}$, formally this assumption can be written: $F_{Y_0|R=0}(Y_0) = F_{Y_{-1}|R=0}(Y_{-1})$ almost surely. Making this assumption, I turn to the union effect on average earnings by pre-election earnings quintile to get a more direct measure of how unionization affects individuals at different points in the distribution. The results suggest that union representation raises the earnings of those at the lower end of the pre-election earnings distribution, and lowers the earnings of those at the upper end. These findings are shown in Figure 2-9, which plots RD estimates of the average treatment effect of unionization on

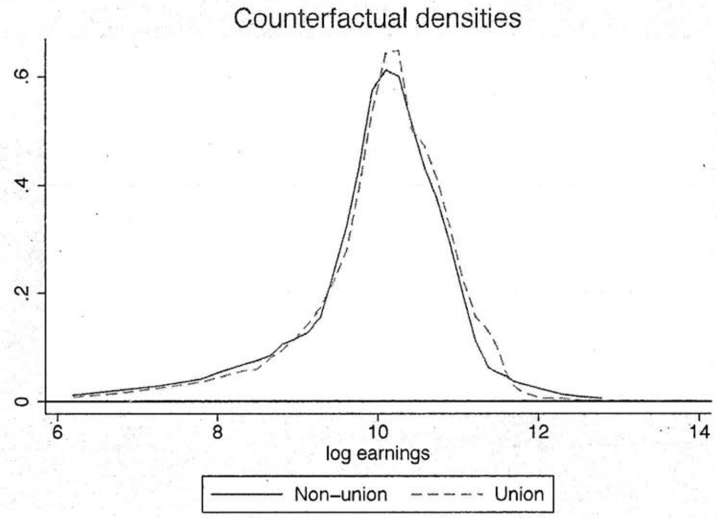

Figure 2-8: Estimated counterfactual densities of log annual earnings conditional on a close union election.

log annual earnings by quintile of pre-election earnings. The estimated effects are between .07 and .09 for the bottom three quintiles, but fall to -.05 for the highest quintile of pre-election earnings. These effects mirror the quantile treatment effects reported above, and suggest that individuals at the lower end of the the earnings distribution prior to a close election can expect to see their earnings increase, while individuals at the upper end may see their earnings decrease as a result of union representation. The similarity of the pattern of effects in Figures 2-6 and 2-9 also suggests that, as a first approximation, the assumption that ranks are preserved across union status may not be unreasonable, and the quantile treatment effects roughly correspond to the expected effect on an individual at a given point in the earnings distribution.

The earnings compression implied by the estimates is consistent with the theoretical predictions in Section 2.2, but does it reflect higher wages for lower-skilled workers, and a reduced return to skill, as the model suggests? Unions may compress earnings for other reasons, including shifting risk to the employer, which could be pareto-improving for workers (Burda, 1995).[8] To get a fuller picture of the welfare consequences of unionization I turn to estimates of the effect on worker retention. The results suggest unionization increases

[8]The literature identifies several other possible sources of union wage compression. Freeman and Medoff (1984) argue that unions provide more efficient levels of public goods in the workplace. Acemoglu and Pischke (1999) show union wage compression can encourage firms to invest in employees' general human capital. Freeman (2005) and Budd (2004) discuss non-wage benefits.

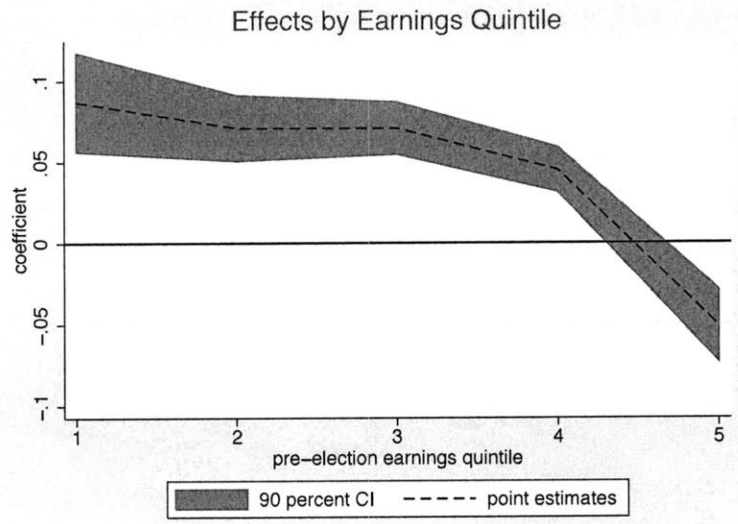

Figure 2-9: Estimates and 90-percent confidence intervals for the average effect of unionization on log annual earnings, by quintile of pre-election earnings.

retention among workers at the lower end of the pre-election earnings distribution, while increasing turnover among workers at the higher end. These findings can be seen in Figure 2-10, which plots estimates and confidence intervals for the effect of unionization on an indicator for retention 10 quarters after the election, by quintile of pre-election earnings. The figure shows unionization has a significantly positive effect on retention for the bottom two quintiles, essentially no effect for the middle quintile, and a significantly negative effect for the top two quintiles. The effect ranges from around 5 percent for the bottom quintile to negative 10 percent for the top quintile, with an overall effect of -.043 (s.e. = .004). The pattern of effects on retention supports the view that unionization makes employment differentially more attractive for lower earners relative to higher earners, consistent with the model's interpretation that union wage compression reflects higher pay for lower-skilled workers and a reduced return to skill.

The effect on retention in Figure 2-10 also has implications for selection into employment at a unionized establishment. Card (1996) developed a two-sided selection model incorporating both employer and employee behavior. If unions compress the distribution of wages, employers are more likely to want to retain (or hire) high-skilled workers, while lower-skilled workers are more likely to want to stay. Figure 2-10 suggests that on net, selection on the part of employees dominates.

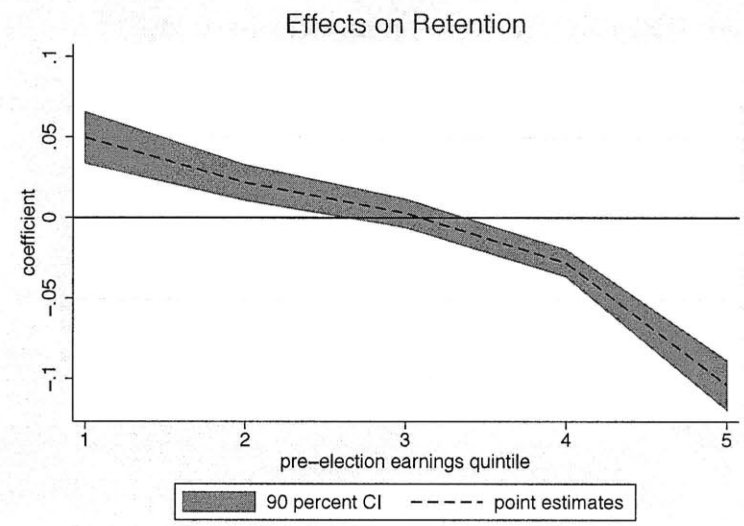

Figure 2-10: Point estimates and 90-percent confidence intervals for the average effect of unionization on retention 10 quarters after the union election, by quintile of pre-election earnings.

The results in this section provide evidence for the union wage compression found in previous regression-based studies, but against a large average effect of unionization. Thus the principal substantive conclusion of Freeman (1993), Card (1996), DiNardo, Fortin, and Lemieux (1996) that unions reduce dispersion holds up even in a quasi-experimental setting, but the finding in these studies and many others that union earnings are on average significantly higher than non-union earnings appears to be largely a selection effect, consistent with DiNardo and Lee (2004), at least for the typical case of a close union election. When unions win by a wider margin, Lee and Mas (2008) find large negative effects on employer profitability, suggesting that the effect on average earnings may increase with the union's vote share. Interpreted in the light of the model in Section 2.2, plants where the union barely wins may have fewer sunk costs (that is, lower ϕ_K), and thus unions extract less from the employer and garner fewer votes compared to plants where ϕ_K is large. Empirical studies looking directly at sunk costs and quasi-rents have also found that unions have a larger effect on employers and industries where sunk costs and rents are higher (see, e.g., Freeman, 1983).

2.6 Summary and Conclusions

Quasi-experimental estimates based on close union certification elections show unionization substantially compresses the distribution of employee earnings. Union representation raised the tenth percentile of earnings by about 25 log points with a much smaller effect in the middle of the distribution, a large negative effect on the upper tail of earnings, and little effect on the average. Estimates of the union effect on employee retention by quantile of pre-election earnings showed a similar pattern: among workers in the bottom two quintiles, unionization significantly increased retention, it had little effect on retention in the middle quantile, and significantly decreased retention in the top two quantiles.

The pattern of effects on the distribution of earnings and worker retention is consistent with a model where unions pursue a wage schedule to achieve political objectives. A union whose growth as an institution depends on new unionization has incentives to set wage schedules to maximize the probability of winning certification elections. The theoretical model in the paper showed unions will raise the wages of lower-skilled workers, but reduce the return to skill, resulting in a compressing effect on the distribution of workers' earnings. The empirical results on worker turnover by earnings quantile also support this interpretation. Further estimation and testing of the model is a subject for future research.

The results imply that unions close to the margin of victory unambiguously reduce dispersion in the overall earnings distribution, since they compress earnings within the unionized sector, but have little effect on the average union earnings gap. Deunionization therefore explains part of the increased inequality in the U.S. income distribution since the 1970s, as Freeman (1993), Card (1996), and DiNardo, Fortin, and Lemieux (1996) also found. A crude estimate of how much of the increased inequality the fall in unionization rates can explain may be obtained using the sample of full-time, full-year private sector wage and salary workers from the 1979 and 2009 Current Population Survey (King, Ruggles, Alexander, Leicach, and Sobek, 2009). The variance of log earnings in this sample increased by about 26 percent from 1979 to 2009, while the private sector unionization rate fell from about 25 percent to 8 percent (Bureau of Labor Statistics, 2010; DiNardo, Fortin, and Lemieux, 1996). Assuming the deunionization occurred among workers at marginally unionized plants, deunionization accounts for about 13.5 percent of the increase in the variance of log earnings.[9] This estimate is close to the 15-20 percent found by Card (2001) and the 6-21 percent found by DiNardo, Fortin, and Lemieux (1996).

The estimates apply only to workers at private sector establishments where a close union election was held. While the estimates thus reflect the causal effect of typical private sector unionization in recent years, they miss the effect of public sector unionization, which now accounts for the majority of U.S. union membership (Bureau of Labor Statistics, 2010). More research is needed on the effects of public sector unionization and on the mechanisms driving those effects.

[9] 13.5% was arrived at as follows. The distribution of unionized potential log earnings in Figure 2-8 has a variance .087 less than the non-unionized distribution. The difference in unionization rates from 1979 to 2009 is $25\% - 8\% = 17\%$. Since the means of the unionized and non-unionized potential log earnings distributions are essentially equal, the increase in overall variance due to deunionization is therefore $17\% \times .087 = .0149$, which is 13.5% of the total increase in variance.

References

Abowd, John M., John Haltiwanger, and Julia Lane (2004): "Integrated Longitudinal Employee-Employer Data for the United States," *American Economic Review Papers and Proceedings*, 94, 224–229.

Abowd, John M., Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock (2009): "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J.B. Jensen, and M.J. Roberts, pp. 149–230. University of Chicago Press for the National Bureau of Economic Research, Chicago.

Acemoglu, Daron, Phillipe Aghion, and Gianluca Violante (2002): "Technical Change, Deunionization, and Inequality," in *Carnegie-Rochester Conference Series on Public Policy.*

Acemoglu, Daron, and Jörn-Steffen Pischke (1999): "The Structure of Wages and Investment in General Training," *The Journal of Political Economy*, 107, 539–572.

Ashenfelter, Orley, and George E. Johnson (1969): "Bargaining Theory, Trade Unions, and Industrial Strike Activity," *The American Economic Review*, 59, 35–49.

Atherton, Wallace N. (1973): *Theory of Union Bargaining Goals.* Princeton University Press, Princeton, New Jersey.

Berkowitz, Monroe (1954): "The Economics of Trade Union Organization and Administration," *Industrial and Labor Relations Review*, 7, 575–592.

BLANCHFLOWER, DAVID, AND ALEX BRYSON (2003): "What effect do unions have on wages now and would 'What do unions do?' be surprised?," NBER Working Paper No. 9973.

BLAU, FRANCINE D., AND LAWRENCE M. KAHN (1996): "International Differences in Male Wage Inequality: Institutions versus Market Forces," *The Journal of Political Economy*, 104, 791–837.

BOOTH, ALISON L. (1995): *The Economics of the Trade Unions.* Cambridge University Press.

BRUDNEY, JAMES J. (2005): "Contractual Approaches to Labor Organizing: Supplanting the Election Paradigm?," in *Labor and Employment Relations Association Series, Proceeding of the 57th Annual Meeting*, pp. 106–113.

BUDD, JOHN W. (2004): "Non-Wage Forms of Compensation," *Journal of Labor Research*, 25.

BURDA, MICHAEL C. (1995): "Unions and Wage Insurance," CEPR Discussion Paper No. 1232. London, Centre for Economic Policy Research. Available at http://www.cepr.org/pubs/dps/DP1232.asp.

BUREAU OF LABOR STATISTICS (2010): "Union Members Summary," in BLS annual union membership report, 2009.

CABALLERO, RICARDO J., AND MOHAMAD L. HAMMOUR (1998): "The Macroeconomics of Specificity," *The Journal of Political Economy*, 106, 724–767.

CAMPBELL, DONALD THOMAS (1969): "Reforms as Experiments," *American Psychologist*, 24, 409–429.

CARD, DAVID (1996): "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis," *Econometrica*, 64, 957–979.

——— (2001): "The Effect of Unions on Wage Inequality in the U.S. Labor Market," *Industrial and Labor Relations Review*, 54, 296–315.

CHAMBERLAIN, GARY (1994): "Quantile Regression, Censoring and the Structure of Wages," in *Proceedings of the Sixth World Cogress of the Econometrics Society,* *Barcelona, Spain*, ed. by Christopher A. Sims, and Jean-Jacques Laffont, pp. 179–209. Cambridge University Press, New York.

CHERNOZHUKOV, VICTOR, AND CHRISTIAN HANSEN (2005): "An IV Model of Quantile Treatment Effects," *Econometrica*, 73, 245–261.

DINARDO, JOHN, NICOLE M. FORTIN, AND THOMAS LEMIEUX (1996): "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach," *Econometrica*, 64, 1001–1044.

DINARDO, JOHN, AND DAVID S. LEE (2004): "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001," *Quarterly Journal of Economics*, 119, 1383–1441.

FARBER, HENRY S. (1978): "Individual Preferences and Union Wage Determination: The Case of the United Mine Workers," *The Journal of Political Economy*, 86, 923–942.

——— (1999): "Union Success in Representation Elections: Why Does Unit Size Matter?," Working Paper No. 420, Industrial Relations Section, Princeton University.

FERGUSON, JOHN-PAUL (2008): "The Eyes of the Needles: A Sequential Model of Union Organizing Drives, 1999-2004," *Industrial & Labor Relations Review*, 62, 3–21.

FRANDSEN, BRIGHAM R. (2008): "Nonparametric Estimation of Local Quantile Treatment Effects in the Regression Discontinuity Design," Available at SSRN: http://ssrn.com/abstract=1216062.

FREEMAN, RICHARD B. (1983): "Unionism, Price-Cost Margins, and the Return to Capital," NBER Working Paper No. 1164.

——— (1993): "How Much Has De-unionisation Contributed to the Rise in Male Earnings Inequality?," in *Uneven Tides: Rising Inequality in America*, ed. by Sheldon Danziger, and Peter Gottschalk, pp. 133–163. Russell Sage Foundation, New York.

——— (2005): "What Do Unions Do? The 2004 M-Brane Stringtwister Edition," NBER Working Paper No. 11410.

FREEMAN, RICHARD B., AND JAMES L. MEDOFF (1984): *What Do Unions Do?* Basic Books, New York.

GELFAND, I. M., AND S. V. FOMIN (1963): *Calculus of Variations.* Prentice-Hall, New Jersey, Edited by Richard A. Silverman.

HECKMAN, JAMES J., JEFFREY SMITH, AND NANCY CLEMENTS (1997): "Making the Most Out of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *The Review of Economic Studies*, 64, 487–535.

JUHN, CHINHUI, KEVIN M. MURPHY, AND BROOKS PIERCE (1993): "Wage Inequality and the Rise in Returns to Skill," *The Journal of Political Economy*, 101, 410–442.

KATZ, LAWRENCE F., AND KEVIN M. MURPHY (1992): "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107, 35–78.

KING, MIRIAM, STEVEN RUGGLES, TRENT ALEXANDER, DONNA LEICACH, AND MATTHEW SOBEK (2009): "Integrated Public Use Microdata Series, Current Population Survey: Version 2.0," Machine-readable database, Minneapolis, MN: Minnesota Population Center [producer and distributor].

KREMER, MICHAEL, AND BENJAMIN A. OLKEN (2009): "A Biological Model of Unions," *American Economic Journal: Applied Economics*, 1, 150–175.

LANE, JULIA (2008): "Linking Administrative and Survey Data," Society of Labor Economists 2008 Annual Meeting Presentation.

LEE, DAVID, AND ALEXANDRE MAS (2008): "The Long-Run Costs of Unionization: New Evidence from Financial Markets, 1961-1999," Unpublished manuscript.

——— (2009): "Long-Run Impacts of Unions on Firms: New Evidence from Financial Markets, 1961-1999," Unpublished manuscript.

LEE, DAVID S. (2008): "Randomized experiments from non-random selection in U.S. House elections," *Journal of Econometrics*, 142, 675 – 697, The regression discontinuity design: Theory and applications.

LEWIS, H. GREGG (1986): *Union Relative Wage Effects A Survey.* University of Chicago Press.

McKinney, Kevin L., and Lars Vilhuber (2008): "LEHD Infrastructure files in the Census RDC - Overview Revision 219," U.S. Census Bureau Technical Paper.

National Labor Relations Board (2010): "Procedures Guide," Available on the world wide web at http://www.nlrb.gov/publications/Procedures_Guide.htm.

Ross, Arthur Max (1948): *Trade Union Wage Policy.* University of California Press, Berkeley.

Thistlethwaite, Donald L., and Donald T. Campbell (1960): "Regression-Discontinuity Analysis: An Alternative to the Ex Post Facto Experiment," *Journal of Educational Psychology*, 51, 309–317.

Appendix A: Theory

The optimal wage schedule is part of a subgame perfect equilibrium: it maximizes the union's probability of winning the certification election, given that workers vote sincerely conditional on the wage schedule, and given that the employer invests and hires to maximize profits conditional on the wage schedule and the outcome of the election. In this setting maximizing the probability of winning is equivalent to maximizing the vote share. To see this, consider an election at a plant with N workers voting. Let $V_i(\theta)$ be the i-th worker's voting rule, as a function of the union's choice parameter, θ. The union solves $\max_\theta \Pr\left(\sum_{i=1}^N V_i(\theta) > N/2\right)$. The number of votes is a Binomial random variable with parameters $(P(\theta), N)$, where $P(\theta)$ is the expected vote share. Denoting the cdf of this binomial random variable by F_B, the union's problem can be rewritten as $\max_\theta 1 - F_B(N/2; (P(\theta), N))$. The first order condition is $\frac{\partial F_B}{\partial P}\frac{\partial P}{\partial \theta} = 0$. Since for the Binomial distribution $\frac{\partial F_B}{\partial P} \neq 0$, this reduces to $\frac{\partial P}{\partial \theta} = 0$, which is the first order condition for maximizing the expected vote share.

Intuitively, for a proposed wage schedule to garner a positive vote share in equilibrium, it must result in K earning at least $(1 - \phi_K)(y - v)$. Otherwise, the employer would simply shut down, and all workers would incur the switching cost to find a job elsewhere, and thus would prefer no union in the first place.[10] Thus, although the union would like to set the wage schedule so as to garner the most votes possible, it must also take into consideration

[10] A recent paper by Kremer and Olken (2009) highlights unions' incentives to take into account employers' profitability.

the direct effect of the wage schedule on the firm's profits via the payroll, as well as the indirect effect via the wage schedule's effect on the distribution of workers' human capital employed at the firm. In consequence, any incentive to redistribute rents among workers of different levels of human capital is tempered by the tendency of workers who are losers under the union wage schedule to quit, further tightening the firm's profit constraint. The following proposition formally characterizes the optimal union wage schedule.

Proposition 5 *A subgame perfect Nash equilibrium union wage schedule is* $w^U(H) = (v + r(H, \lambda^*)) H$*, where the union rent schedule* $r(H, \lambda)$ *satisfies*

$$\frac{f_\eta(r)}{F_{\eta-\varepsilon}(r)} = \lambda H \left(1 - (\phi_K(y-v) - r)\frac{f_{\eta-\varepsilon}(r)}{F_{\eta-\varepsilon}(r)}\right),$$

and λ^* *satisfies*

$$\int (\phi_K(y-v) - r(h, \lambda^*)) h F_{\eta-\varepsilon}(r(h, \lambda^*)) \, dF_H(h) = 0,$$

and $F_{\eta-\varepsilon}$ *denotes the distribution function of the random variable* $\eta - \varepsilon$*, the unionization cost net of the switching cost.*

Proof. Working backwards, consider the firm's investment and hiring decision given $r(H) = w(H)/H - v$, the outcome of the election, and workers' quitting decisions. First take the case where the union loses. Then the price of H hasn't changed, and the competitive equilibrium remains optimal for the employer. If any workers quit, the employer hires from the pool of applicants (in this case identical in distribution to the population of workers) to replace them and production continues. If no workers quit, no additional hiring or investment takes place and production resumes as before. Now take the case where the union wins. Since in the short run the production technology takes fixed proportions of H and K, if no workers quit, the available actions are hiring additional H and K, or releasing currently employed H and K. The employer would hire additional H and K only if the return to the additional K is greater than $y - v$ (the purchase price of K). However, the employer must pay at least v in order to hire more H, so the equilibrium cannot involve the employer hiring additional H or investing in more K. Still considering the case where the union wins but no workers quit, the employer releases currently employed H and K if the return to K under the union wage schedule is less than $(1 - \phi_K)(y - v)$. In consequence of constant returns to scale and fixed proportions, if it's optimal to release any H and K, it is optimal to shut down completely. Finally, in the case where the union wins and some workers quit, it is never optimal for the employer to replace the lost workers, for the same reason the employer doesn't hire additional workers. Instead the employer will divest the freed up K (recouping $(1-\phi_K)(y-v)$ per unit). The employer will either resume production with the remaining H and K, or shut down, again depending on whether the return to the remaining K under the union wage schedule is greater than or less than $(1-\phi_K)(y-v)$. In summary, the employer's equilibrium strategy is to divest any K freed up by quitting workers, and continue production with what remains if the return to K is at least $(1-\phi_K)(y-v)$, and shut down otherwise.

Turn now to the worker's decision to stay or leave conditional on $r(H)$ and given the employer's equilibrium strategy. A worker of human capital H chooses to stay if her union rent, $r(H)$, exceeds her cost of unionization net of her switching cost. Thus the equilibrium decision rule is $1(\eta-\varepsilon \leq r(H))$. The density of human capital conditional on staying at the firm is therefore

$$f_{H|\text{stay}}(h)=\frac{F_{\eta-\varepsilon}(r(h))\,f_H(h)}{\int F_{\eta-\varepsilon}(r(s))\,dF_H(s)},$$

where $F_{\eta-\varepsilon}$ denotes the distribution function of the random variable $\eta-\varepsilon$, the cost of unionization net of the switching cost.

Next, consider the workers' voting choice conditional on $r(H)$ and given the employer's equilibrium strategy. No worker will vote for the union if $r(H)$ is such that the employer shuts down production, since all workers would then incur the switching cost ε, and earn the same wage, v, elsewhere. Conditional on $r(H)$ not resulting in a shut-down, a worker of human capital H votes for the union if her rent under the union schedule exceeds her cost of unionization, η. The equilibrium voting function is therefore $1(\eta \leq r(H))$, again subject to the condition that the plant stays open under the wage schedule.

Finally, given the workers' and the employer's equilibrium strategies, the union chooses $r(H)$ to maximize the expected vote share, which is equivalent to maximizing the expectation of $1(\eta \leq r(H))$, subject to the firm earning an ex post return of at least $(1-\phi_K)(y-v)$.

The problem the union solves can be written:

$$\max_{r(h)} \int F_\eta (r(h))\, dF_H(h)$$
$$\text{s.t.} \quad \int (\phi_K (y - v) - r(h))\, h F_{\eta-\varepsilon}(r(h))\, dF_H(h) \geq 0. \tag{2.4}$$

This is a straightforward calculus of variations problem of the type treated by, say, Theorem 1 in Gelfand and Fomin (1963, p. 43). The optimal rent schedule therefore satisfies

$$\frac{f_\eta(r)}{F_{\eta-\varepsilon}(r)} = \lambda H \left(1 - (\phi_K(y-v) - r)\frac{f_{\eta-\varepsilon}(r)}{F_{\eta-\varepsilon}(r)}\right),$$

where λ is the Lagrange multiplier on the minimum profit constraint for the firm. This condition can be solved for $r(H, \lambda)$. Substituting this into the profit constraint, (2.4), and solving for λ gives λ^*, and thus the optimal wage schedule can be written

$$w^U(H) = (r(H, \lambda^*) + v)\, H.$$

■

Appendix B: Data

Construction of the dataset

As described in the text, the dataset used in this paper consists of NLRB certification election results matched to employer-employee wage data from the Census Bureau's LEHD program.

The union certification election records were collected by the NLRB, and in large part maintained by the AFL-CIO. John-Paul Ferguson obtained the data by filing Freedom of Information Act requests with the NLRB, and has made them available for this research. The complete data set covers the period 1963-2006, and contains records from about 250,000 union elections, although the main sample used in the analysis covers the years 1992-2001, including 37,354 elections. The raw data contains results from elections stemming from several different type of petitions, including cases where a union seeks to be certified (RC), an employer seeks an election to remove an existing union (RM), or employees seek to remove a union (RD). I restrict to RC-cases, where a union seeks certification. The dataset contains many duplicate records. In some cases they are true duplicates: one election generated multiple records in the database. In these cases I simply delete the redundant entries. In other cases, multiple entries arise from more than one union being on the ballot. In these cases the relevant union vote share is the largest one; I therefore retain the entry with the largest vote share, and delete the others. Finally, in some cases multiple elections were held at the same establishment because, for example, different groups of workers constituted different bargaining units. Since I can't distinguish between workers in different bargaining units, the relevant vote share is the largest, so again I keep only the entry corresponding to the election where the union received the highest vote share.

The second data component consists of the Employment History Files (EHF) within the LEHD database. As described in the text, the EHF contains employee, employer, and earnings data for each employment relationship that generated at least one dollar of wages. The EHF includes a state employer identification number (SEIN) with each record, and in some cases an identifer for the establishment within the employer, which is important for multi-unit employers. For the cases where there is no establishment identifier, the LEHD provides a Unit-to-Worker (U2W) imputation to assign workers to establishments. The employer name and address of these establishments—obtained from the Business Register's Standard Statistical Establishment List (SSEL)—are then used to link to the union election dataset to determine union coverage status.

The matching process to combine these two data sources is as follows. First, employer name and address information from both the NLRB dataset and the Census Bureau's Business Register (BR) were cleaned and standardized using the SAS Data Quality Server standardization functions. NLRB election records were then matched to BR records by several combinations of state, county, city, employer name, street address, and industry code. The match was performed iteratively in descending order of strictness. The cutoff level of strictness was determined by hand checking matches from each iteration, and stopping once match quality dipped below 95 percent. The matched BR records were then linked to employers in the LEHD's Employer Characteristics File (ECF) by the Business Register Bridge (BRB) via state, year, county, Employer Identification Number (EIN) and two-digit industry code. Finally the work histories (including earnings) of all individuals employed at the matched employers during the quarter of the certification election were drawn from the Personal History File (PHF), using the Unit-2-Worker imputation to complete the match in the case of multi-unit employers.

Defining the Running Variable

A critical feature of the regression discontinuity design is the running variable (in this case the amount by which the union's vote share exceeds 50 percent). As DiNardo and Lee (2004) point out, care must be taken when defining this variable to avoid biasing results toward the smallest employers. I follow their procedure of first subtracting .5/(# votes cast) from the vote shares where an even number of votes were cast, and then binning the resulting modified shares so that all elections with a share between .50 and .55 are assigned .525, and so forth. Finally, only elections where the number of total votes cast exceeded 10 were kept in the analysis.

Sample Selection

The sample included in the main analysis consists of those workers who have non-missing wage data for the one-year period beginning two quarters after the closing date of the election. I also condition (approximately) on full-time/full-year workers by keeping only those workers whose wages in the year prior to the election exceeded (in 2000 dollars) 20 hours/week $\times$ 40 weeks/year $\times$ \$5/hour = \$4000/year. Although crude, this conservative approximation to full-time, full-year status is based on pre-determined wages and so does not affect the validity of the estimation, and aids in interpretation.

Chapter 3

Did Vietnam Veterans Get Sicker in the 1990s? The Complicated Effects of Military Service on Self-reported Health[1]

3.1 Introduction

The difficulties faced by many Afghanistan and Iraq war veterans have once again drawn attention to the fact that military service can have long-term health consequences. Care for injured and disabled veterans imposes a burden on soldiers, their families, and, in a less personal but still important way, on the government agencies that provide health care and disability insurance to veterans. These social insurance systems support almost three million sick and disabled veterans. Veterans Administration (VA) support programs increasingly serve a relatively young population made up of veterans of post-Korea conflicts. Vietnam veterans constitute the largest group receiving veterans disability compensation (VDC), with almost one million beneficiaries, about one third of the total. At the end of fiscal year 2006, the population receiving VDC amounted to roughly 12 percent of Vietnam-era veterans and 15 percent of Gulf War veterans, exceeding VDC take-up rates of 5 percent among Korean-era veterans and 10 percent among veterans of WWII. Moreover, payments

[1]This chapter is joint work with Joshua Angrist and Stacey Chen

to Vietnam veterans have expanded to almost half of VDC costs, partly because Vietnam veterans are disproportionately likely to receive the maximum payment allowed (Veterans Benefits Administration, 2007).

The most visible health concerns for veterans are the long-term consequences of combat injury. Battlefield injuries can be individually devastating and socially costly for years after a conflict ends. Fortunately, acute injuries are less common among veterans of recent conflicts than they were in WWII (U.S. Bureau of the Census, 2006). At the same time, an increasing fraction of veteran disability claims in the past two decades has been for chronic conditions that were not necessarily apparent on the battlefield. These conditions include post-traumatic stress disorder (PTSD), hearing loss, and diabetes. Evidence for the importance of PTSD among Vietnam veterans comes in part from the pioneering draft-lottery study by Hearst, Newman, and Hulley (1986), which showed elevated civilian suicide rates for draft-eligible men. Among Gulf War veterans, a large and growing health concern stems from a collection of symptoms with no specific identifiable cause known as Gulf War syndrome. The question of whether military service is indeed the root cause of these symptoms continues to be debated, but they are usually presumed to be service-connected and therefore covered by VDC.[2]

The civilian re-entry experiences of each veteran cohort are in many ways unique, but there are some striking similarities. The debate over Gulf War syndrome echoes a similar controversy surrounding the rise in disability claims by Vietnam veterans—a rise that accelerated in the late 1990s and continues today. Until very recently, claims by Vietnam veterans were the source of most VDC claims growth. After 2002, this growth is partly attributable to the Veterans Benefits Administration's designation of diabetes as a service-related disability linked to the herbicide Agent Orange (Autor and Duggan, 2008).[3] Perhaps surprisingly, however, much recent claims growth is due to new PTSD claims by Vietnam veterans (Rosenheck and Fontana, 2007). The recent growth in Vietnam veterans' disability claims makes the long-term health consequences of Vietnam-era service an important contemporary policy concern.

An assessment of the link between Vietnam-era military service and long-term disability

[2]Key studies of this question include Research Advisory Committe on Gulf War Veterans' Illnesses (2008) and Medical Research Council (2003). See Iversen, Chalder, and Wessely (2007) for a review.

[3]The Veterans Benefits Administration is the part of the VA that administers VDC and other benefit programs.

also contributes to a broader understanding of the likely health and social insurance costs of other wars. The Vietnam War lasted longer and was much more costly in terms of fatality and injury rates than more recent conflicts. Consistent with this, Vietnam veterans are much more likely to receive the maximum VDC benefit than any other cohort. We might therefore expect the health and social insurance consequences of Vietnam-era service to provide a rough upper bound on the long-term health consequences of military service for veterans of recent wars. The experience of veterans from earlier wars provides the best available evidence on the likely consequences of military service for more recent veterans, a point made in a recent assessment of the long-run costs of the Iraq conflict by Stiglitz and Bilmes (2008).

The purpose of this paper is to provide new evidence on the long-term health impact of Vietnam-era military service as the affected cohorts reached their late 40s and early 50s. Because employment is closely associated with health, we also look at veterans' labor force status. To solve the problem of selection bias inherent in comparisons of outcomes between veterans and non-veterans, we use the draft lottery to construct instrumental variables for Vietnam-era service.[4] Our empirical strategy relies on the 1-in-6 sample of the 2000 U.S. Decennial Census. The 2000 Census provides an exceptionally large sample and, uniquely among large representative samples, contains the birthday information required to determine draft lottery numbers. Moreover, in addition to the usual labor force status variables, the 2000 Census long form asks respondents about disabilities along a variety of dimensions, with a distinct category for disabilities that affect work. Our results show no overall causal effect of Vietnam-era veteran status on employment, labor force participation, or work-limiting disabilities (that is, long-lasting physical or mental health conditions causing difficulty working). On the other hand, we find a large increase in federal transfer income and modest effects on disabilities that census respondents describe as not limiting work.

An important feature of our analysis is an exploration of veteran effects that vary with veterans' predicted wages and schooling. High replacement rates (i.e., the ratio of disability income to prior earnings) have made Social Security Disability Insurance (SSDI) and Supplemental Security Income (SSI) an increasingly attractive alternative to employment for low-skilled men not yet old enough to retire, as Black, Daniel, and Sanders (2002), Autor

[4]Earlier work by Angrist (1990) and Angrist and Chen (2008) uses the draft lottery to evaluate the earnings consequences of military service.

and Duggan (2003), and Duggan, Singleton, and Song (2007) have argued. Motivated by the possibility of similar interactions with VDC, we look for differences in the impact of veteran status across quantiles of predicted wage levels and schooling. The results of this investigation show a strong interaction: in contrast to the small overall effects, estimates for veterans with low predicted earnings show a large negative effect on employment and a marked increase in disability rates (again, mostly for disabilities not limiting work). Moreover, there is little evidence that this variation can be explained by variation across skill levels in the likelihood of serving in combat or a war zone. As measured in the 1987 Survey of Veterans, the likelihood of a service-connected disability and combat exposure is similar for high school graduates and dropouts. These results therefore suggest that the causal effects of Vietnam-era service on employment and disability for less educated veterans reflect something other than wartime injuries. A leading alternative explanation is the relative attractiveness of VDC for less-skilled men and the work disincentives embedded in the VDC system.

The paper is organized as follows. The next section uses statistics on VDC and data from the CPS to identify the primary health concerns of Vietnam veterans and to describe recent changes in veterans' disability status, beneficiary status, and employment rates. Section 3.3 discusses the descriptive statistics and first-stage estimates from the 2000 Census that provide a foundation for the draft-lottery-based causal analysis in Section 3.4. Section 3.3 also briefly discusses the impact of Vietnam veteran status on mortality, since this is a possible source of selection bias in our analysis. Section 3.4 reports overall disability and employment effects and effects by predicted wage and schooling group. These results show important differences in effects across skill groups. Section 3.5 discusses the link between schooling and variables related to combat or war-theatre exposure and interprets the other findings in the paper. Finally, Section 3.6 concludes.

3.2 VDC and Health in the Vietnam Cohort

Veterans disability compensation (VDC) increases with the recipient's combined disability rating (CDR), which is the aggregate of ratings for all diagnoses for which VDC is awarded. Veterans with a zero percent CDR get no monthly payment but are eligible to use the VA health care system. The largest awards go to veterans deemed to be 100 percent disabled.

Veterans with a single disability rated at 60 percent or more, or a combined rating of 70 percent or more plus a single disability rated at 40 percent or more, can receive an Individual Unemployability (IU) benefit if the Veterans Benefits Administration determines that they cannot work by virtue of their disabilities. An IU determination generates payments at the 100 percent CDR level. As noted by Autor and Duggan (2008), the IU contingency generates a substantial implicit tax on the earnings of VDC awardees.[5]

Our descriptive analysis focuses initially on VDC award and diagnosis data for Vietnam veterans in the period 1999-2005 because this is a time when the Vietnam-era VDC beneficiary population was growing and changing in important ways, as shown in Table 1, which list the nine most common diagnoses among Vietnam-era veterans (as determined in 2005) and the number of disability recipients under these diagnoses for the years 1999, 2001, 2003, and 2005. The top panel shows that in fiscal year 2005, over 900,000 Vietnam veterans were receiving VDC, up sharply from about 736,000 in 1999 (this can be compared to a Vietnam-era population of 8 million veterans, of whom about 3 million served in Vietnam or nearby). The number of disabling conditions per Vietnam-era recipient also increased, from 2.76 to about 3. Panel C shows the same trend is true for veterans of more recent conflicts, such as the Gulf War. In contrast, statistics in Panel B for Korean-era cohorts show lower proportions receiving compensation, less growth in the numbers receiving compensation, and fewer disabilities per veteran.

In fiscal year 2005, the most widely compensated disability among Vietnam veterans was diabetes, with more than 190,000 recipients. This can be seen in the bottom of Panel A in Table 1. Diabetes was recognized as a service-related disability beginning in fiscal year 2002, in response to evidence of a possible link with exposure to the Agent Orange herbicide used by US forces during the Vietnam War. The growth in diabetes claims from zero in 2001 to many thousands in 2003 is not a result of new cases of diabetes, but rather reflects the fact that diabetes was a newly recognized service-related condition.

The most prevalent condition for which Vietnam veterans received compensation from 1999 to 2003 is PTSD, with about 91,000 claimants in 1999 and 143,000 in 2003. PTSD

[5]In fiscal 2004, VDC award amounts ranged from about $100 per month for veterans with a 10 percent CDR to almost $2,300 per month for veterans with a 100 percent CDR (VA Office of the Inspector General, 2005). There is a much larger increase in benefits in the step from 90 to 100 percent CDR than at other steps. VDC benefits are not subject to federal income or payroll taxes and have usually kept pace with inflation. Appendix II of U.S. General Accountability Office (2006) estimates that an IU determinations adds about $348,000 to the lifetime present value of VDC payments for a 45 year old man with a 60 percent schedular CDR.

has long been a health concern for Vietnam veterans; the incidence of PTSD among Vietnam-era veterans is much larger than that for veterans of the Korean War, the Gulf War, and WWII (from statistics not shown here). Perhaps surprisingly, the number of Vietnam-era PTSD claimants doubled between 1999 and 2005, long after the Vietnam war ended. Although the recent increase in PTSD claims is sometimes attributed to the psychological impact of post-September 11th conflicts, Table 1 shows a marked increase between 1999 and 2001 (data for 2001 are from July), and a dramatic jump between fiscal 2001 and fiscal 2003, before the wars in Afghanistan and Iraq had begun to generate large numbers of casualties. Moreover, an analysis of veterans' use of PTSD treatment services in the 6 months before and after the September 11 attacks failed to uncover a short-term increase in the number of veterans seeking treatment (Rosenheck and Fontana, 2007). At the same time, psychiatrists and others involved with the treatment of PTSD have noted anomalies that point to financial motives on the part of some treatment-seeking veterans. These include volatile or implausible PTSD symptom descriptions and combat-experience reports, the apparent ineffectiveness of treatment for combat-related PTSD (therapy for non-military PTSD has been shown to be effective), and a review of military records that turned up evidence of combat exposure for only 41 percent of those seeking PTSD treatment at a VA clinic (see Frueh, et al. 2005 for this last finding and a summary of related results).

The increase in PTSD among Vietnam veterans has led to a number of government studies, motivated by the question of whether this increase reflects a true deterioration in health or a change in VDC eligibility screening standards and diagnostic criteria (see, e.g., VA Office of the Inspector General, 2005 and Institute of Medicine and National Research Council, 2007). A related concern is the growing proportion of PTSD beneficiaries designated IU, which increased from 14 percent in 1999 to almost 30 percent in 2006. Over one-third of IU beneficiaries in 2006 had PTSD as either a primary or a secondary diagnosis (U.S. General Accountability Office, 2006).[6]

Evidence for a regulatory or administrative explanation of the growth in the number of Vietnam veterans receiving VDC comes in part from state variation in average VDC payments. Specifically, a major contributor to cross-state differences in VDC awards appears

[6]This statistic is from Tables 5-9 and page 150 in Institute of Medicine and National Research Council (2007). A report by the VA Office of the Inspector General (2005) notes that there is more discretion in IU determinations than in the result of the CDR rating system. An earlier report along these same lines (U.S. General Accounting Office, 1987) recommended that IU determinations be subject to an evaluation by the VA's vocational services division, but this has yet to be implemented.

to be variation in the likelihood that otherwise similar cases are designated 100 percent disabled as a result of IU or PTSD (VA Office of the Inspector General, 2005). Along the same lines, the General Accountability Office found that the number of IU beneficiaries was increased by the fact that, beginning in 1999, the Veterans Benefits Administration no longer required IU recipients to submit any kind of paperwork to maintain their IU status (U.S. General Accountability Office, 2006).[7] Moreover, around this same time, the Veterans Benefits Administration began to *presume* IU eligibility in some cases where veterans would previously have been required to actively file an IU claim (Cooper, 2005).

3.2.1 Disability and Beneficiary Status in the CPS

A longer view of trends in Vietnam veterans' disability status appears in Figure 1, which shows the average amount and incidence of VA-source income. These data come from the Current Population Survey (CPS) and are described in the data appendix. The sample includes Vietnam-era and Korean-era cohorts over the period 1988-2005.[8] Changes in VA-source income are mostly due to changes in VDC, since the Vietnam-era GI Bill expired in 1989. Figure 1 shows a marked increase in the average VA income of Vietnam veterans in the late 1990s, and a sharper contrast in VA income between Vietnam-era and Korean-era veterans after 2000. During the same period, VA income levels were fairly flat for Korean-era veterans. As shown in Panel B, the relative likelihood that a Vietnam veteran received any VA income compared to Korean-era veterans also jumped in the late 1990s, though this series is noisier than the average income series.

The growing share of Vietnam-era veterans receiving VA income in the past two decades comes mainly from growth in the probability of receiving VDC, shown in Panel C of Figure 1, which contrasts the probability of VDC receipt for Vietnam-era and Korean-era veterans over the same period. This figure closely parallels the plot in Panel B.[9] Overall, the VA income trend for Korean-era veterans has been flattening or declining since the late 1990s, in marked contrast to the rapid increase in VA income receipt among Vietnam-era veterans over this period. Although Korean-era veterans are older, VA income receipts for Vietnam veterans by 2005 are at much higher levels than for Korean-era veterans at the beginning

[7]This order was rescinded in 2005 (Philpott, 2005).

[8]The comparison should be interpreted with caution, since Korean-era veterans are of course older than Vietnam veterans. For this reason we also compare Vietnam veterans to non-veterans of the same age in subsequent figures to get a more complete descriptive picture.

[9]Data for the figures in this subsection are described in the appendix.

of the series, when Korean-era veterans were of a similar age.

Direct measures of self-reported disability rates and a measure of poor health, plotted in panels A and B of Figure 2, also increased in the late 1990s, both in absolute terms and relative to the trend among non-veterans. This increase may reflect a deterioration in the health of Vietnam veterans, but the sharpness of the break suggests that policy or regulatory changes may also play a role. Consistent with the regulatory hypothesis, Duggan, Rosenheck, and Singleton (2006) conclude that modest changes in medical eligibility criteria for federal disability programs can substantially affect program enrollment. Regulatory changes may in turn influence self-reports of health if these measures are at least in part endogenous in the sense that they are caused by program use (a point made by Bound and Waidmann (1992) regarding social security disability programs). In the CPS, there is a further mechanical link between disability income and disability assessment since the CPS disability question is a screener for questions about disability income.

As we might expect given the growing importance of IU claims in the overall VDC caseload, Panel A of Figure 3 shows that the employment rate of Vietnam veterans also dipped in the late 1990s, relative to the non-veteran trend. Although employment should fall as the Vietnam cohort ages, the figure shows a dip relative to non-veterans of the same age with the gap by veteran status eventually increasing over time, after a period in which veteran and non-veteran employment rates had moved roughly in parallel. Panel B of Figure 3 shows that this relative decline is associated with a decline in self-reported health: the fraction of Vietnam veterans reporting that they quit a job or retired for health reasons ticked up sharply in 1998, and eventually pulled away from the same measure for non-veterans in the cohort, although the measure is somewhat noisy. Following a brief review of related work, our empirical analysis attempts to determine whether a causal effect of Vietnam-era military service on health can explain the relative deterioration in Vietnam veterans' self-reported disability status and employment rates.

3.2.2 Related Work on Military Service and Health

The question of how military service affects civilian health is of long-standing concern to veterans and policy-makers. As noted in the introduction, one of the most controversial issues in the health arena is the proper clinical response to Gulf War Syndrome (Iversen, Chalder, and Wessely, 2007). The growth in PTSD diagnoses among Vietnam veterans has been similarly controversial (Rosenheck and Fontana, 2007). Perhaps not surprisingly, given the numbers of men involved and the unique features of each era, the subject of military service and health has generated a large literature covering each service era back to WWII. A comprehensive review of these literatures is beyond the scope of our paper, but it' s worth emphasizing the importance of selection bias in this context. This selection problem is highlighted by Seltzer and Jablon (1974), which shows that WWII veterans live longer than non-veterans born in the same years, primarily due to lower death rates from conditions that would have made them ineligible for service.

A number of earlier studies have used instrumental variables in an effort to eliminate selection bias in estimates of the health effects of military service, as we do here. Bedard and Deschenes (2006) used cohort-dummy instruments to show that military service during World War II and the Korean conflict led to higher mortality from smoking-related causes, apparently because soldiers had access to free or subsidized cigarettes. In contrast, using draft lottery instruments, Eisenberg and Rowe (2008) found no evidence of a lasting increase in smoking by Vietnam veterans (who did not get as large a cigarette subsidy as WWII veterans). Hearst, Newman, and Hulley (1986) found excess suicide and motor vehicle death rates among draft-eligible men. Excess deaths from these causes might be due to PTSD. But a re-analysis of the HNH data by Angrist, Imbens, and Rubin (1996) was less conclusive. Similarly, also using draft-lottery instruments, Dobkin and Shabani (2006) found no clear link between Vietnam-era service and a range of health outcomes measured in the National Health Interview Survey.[10]

[10]Other draft lottery studies include Goldberg, Richards, Anderson, and Rodin (1991), who found no evidence of increased alcohol consumption among draft-eligible men and Hearst, Buehler, Newman, and Rutherford (1991), who found no increase in AIDS among draft-eligible men, although many Vietnam veterans stationed overseas were thought to have experimented with intravenous drugs.

3.3 Census Data and the Draft-Lottery First-Stage

3.3.1 The 2000 Census 1-in-6 File

The 2000 long-form census sample includes approximately one-sixth of U.S. households.[11] For the purposes of this study, we created an extract from this sample consisting of U.S.-born men residing in the 50 States and the District of Columbia, born between 1948 and 1952. The cohorts of 19-year-olds at risk of conscription in the draft lotteries were born from 1950-52. Men born in 1948 and 1949, who were also affected by the 1970 lottery, are included as well. The estimation sample contains more than 1.14 million whites and about 155,000 nonwhites.

Roughly 31 percent of men born 1948 to 1952 served in the Vietnam era and about 44 percent were draft-eligible. The average age in the cohort is 49. These and other descriptive statistics appear in Table 2, which reports means by veteran status and race (means for whites appear in Panel A and means for nonwhites appear in Panel B). Many men report having some kind of disability—about 20 percent of whites and a third of nonwhites. Based on a question asking respondents whether they have a disability that causes difficulty working, we categorized disabilities as *work-limiting* or *non-work-limiting.*[12] While this distinction may be imprecise and subjective, it provides a simple measure of severity and may also be informative about the presence of disabilities that might support SSDI or VDC-IU claims. Among white veterans, the regression-adjusted labor force non-participation rate is 3.3 percentage points larger for those with a work-limiting disability than for those with a non-work-limiting disability. Likewise, the regression-adjusted probability of SSDI receipt is 3.4 percentage points larger for those with a work-limiting disability. Both of these differences are significant at conventional levels.[13]

[11]The 1-in-6 long form sample is the basis for the publicly available PUMS files. These files, documented in US Census Bureau (2005), are simple random samples drawn from the 1-in-6 file, though the 1-in-6 file is not a simple random sample from the census sampling frame. Rather, the Census Bureau reduces the sampling rate in more densely populated areas. Adjustment for variation in sampling rates is made here by using the weighting variables that are included in the long-form file. These weights adjust for non-response and for non-random sampling, and are designed to match external population totals by age, race, sex and Hispanic origin. In practice, weighting matters little for our results.

[12]The work-related disability question asks: *Because of a physical, mental, or emotional condition lasting 6 months or more, does this person have any difficulty working at a job or business?* The complete set of 2000 Census disability questions appears in the data appendix.

[13]These estimates come from regressions that include dummies for state of birth, year of birth, and month of birth. Among white veterans, the estimated effect of a non-work-limiting disability (relative to no disability) on labor force non-participation is about 0.27 (s.e.=0.003), and the estimated effect on SSDI receipt is about 0.10 (s.e.=0.002).

White veterans have somewhat higher disability rates than white non-veterans, while disability rates differ little by veteran status for nonwhites. Table 2 also shows that both white and nonwhite veterans are much more likely than non-veterans to report having income in a category that includes VDC. This is coded from an *other income* question that asks about sources of income received regularly such as veterans' (VA) payments, unemployment compensation, child support, or alimony. Since our all-male sample probably has no income from child support or alimony, and employment rates differ little by veteran status for whites and are higher for nonwhite veterans, the other income differential by veteran status is most likely due to VDC.[14] In our sample, about 11 percent of white veterans have other income, a sharp contrast to the 3.7 percent of non-veterans who have this sort of income. Among nonwhites, 14 percent of veterans and 5.2 percent of non-veterans have other income.

Social Security income is measured by two variables in the same panel of Table 2: Supplemental Security Income (SSI), and Social Security income excluding SSI. Since Vietnam-era cohorts are too young to have retired, and are unlikely to qualify for benefits under the means-tested SSI program, their Social Security income is mostly from Social Security Disability Insurance (SSDI). About 4 percent of whites and 6 percent of nonwhites receive SSA income other than SSI. Among whites, the proportions receiving SSI are 1.3 percent for veterans and 1.9 percent for non-veterans. Among nonwhites, veterans are also less likely to receive SSI than non-veterans (2.9 percent compared to 5.1 percent). Finally, we constructed an indicator for men who receive any federal transfers–either other income, Social Security income, or SSI. Not surprisingly, given the other income differential, both white and nonwhite veterans are much more likely to have federal transfer income of some sort.[15]

Table 2 also reports statistics for the specific types of disabilities identified in the census. The incidence of each disability is slightly higher for veterans than for non-veterans

[14]Income from military disability pensions received directly from the defense department is most likely captured by another census variable that asks about retirement, survivor, or disability pensions excluding Social Security. The *other income* variable might also include payments received under the GI Bill, but few veterans in our cohort were likely to have still been active in GI Bill supported training. Vietnam-era GI Bill eligibility expired in 1989. The 2000 Census income questions used in our study appear in the data appendix.

[15]Eligible claimants can receive both VDC and SSDI benefits without any reduction from either program; see Autor and Duggan (2008). In 2005, about 61 percent of VDC claimants with an IU rating received SSDI (Christensen, McMahon, Schaefer, Jaditz, and Harris, 2007). Few veterans receive SSI because SSI is means tested and because most veterans have a work history that qualifies them for the more generous SSDI program.

among whites, while the difference in specific disability rates by veteran status is small for nonwhites. Among veterans, the most commonly reported disability is related to mobility (identified in a question asking about going outside the home alone to shop or visit a doctor' s office). The second most common disability is associated with restricted physical activities such as walking, climbing stairs, reaching, lifting, or carrying. Mental disabilities, the third most common type of impairment, are recorded in response to a question about difficulty learning, remembering, or concentrating.

The bottom of the table shows descriptive statistics for two labor force status variables that might be related to self-reported disability status, *not working* (one minus employment) and *not in the labor force*. White veterans and non-veterans are about equally likely to be working or in the labor force. Among nonwhites, veterans work more.

Finally, since differential effects across skill groups play an important role in our analysis, we also provide summary statistics for veterans and non-veterans by skill group in appendix Table A2.

3.3.2 The Draft-Lottery First Stage

The first draft lottery, held in December 1969, affected men born in 1944-50 who were at risk of conscription in 1970, while subsequent draft lotteries involved 19-year-olds only. Men born in 1951 were at risk of conscription in 1971, and men born in 1952 were at risk of conscription in 1972. Men born in 1953 were assigned lottery numbers in 1972, but there were no draft calls in 1973.

Each lottery was associated with a draft-eligibility ceiling or cut-off. Men with a random sequence number (RSN) below the ceiling were draft-eligible, while men with an RSN above the ceiling were draft-exempt. The draft-eligibility ceiling was 195 in the 1970 lottery, 125 in the 1971 lottery, and 95 in the 1972 lottery. Draft eligibility is highly correlated with Vietnam-era veteran status, but the link is far from deterministic. Many men with draft lottery numbers below the ceiling were able to avoid conscription through an occupational or educational deferment, or because of poor health or low test scores, while many with lottery numbers above the ceiling volunteered for service. Throughout the Vietnam era (1964-1975), most soldiers were volunteers. Using the draft lottery as an identification strategy yields estimates of the effects of military service specific to the set of "draft lottery compliers": those individuals who were or would have been induced to serve by being draft eligible, but who would not have served otherwise. While our results may not apply to volunteers, compliers make up a substantial fraction of our sample, as the first stage estimates we report below indicate.

In the sample of men born 1948-52, the effect of draft eligibility on Vietnam-era veteran status is .112 for whites and .072 for nonwhites, as shown in panels A and B of Table 3. Draft-eligibility effects for men born 1944-47 (not reported here) are small so we omit these cohorts.

Our primary IV strategy uses a draft-eligibility dummy as an instrument for veteran status. However, in an effort to produce more efficient 2SLS estimates by exploiting within-eligibility changes in the probability of enlistment, we also work with an instrument set constructed from five lottery-number groups and interactions of these groups with year of birth (an instrument set we call $5zx$). The five lottery-number groups are constructed using RSN cutoffs of 95, 125 and 195, and two intermediate points, 160 and 230. The reference group consists of those with RSNs above 230. The intermediate points capture the small but significant increases in the probability of service for RSNs just above the cut-offs.

The first column in each panel of Table 3 reports estimates of the lottery-group first stage in pooled samples.[16] For example, column (1) shows that men born 1948-52 with RSNs up to 95 were .128 more likely to serve than men with RSNs above 230. The next group, with RSNs 96-125, was .082 more likely to serve than the reference group; the next group was .058 more likely to serve; the next group after that was .044 more likely to serve; and the last group with RSNs 196-230 was .006 more likely to serve, although they were above the cut-off at 195. All of these first-stage effects are precisely estimated and significantly different from zero. The draft lottery may have induced some with RSNs above the cutoff to volunteer to obtain better terms of service (e.g., choice of branch of service) in case they were to be drafted eventually. As with the draft-eligibility effects, estimates of lottery group effects are consistently smaller for nonwhites than for whites. F-statistics in the pooled 1948-52 sample range from 134 for nonwhites to nearly 2300 for whites. The $5zx$ first stage appears in columns (2) through (6) of Table 3. Partial F-statistics for the marginal contribution of $5zx$ in a model that includes lottery-group main effects are on the

[16]The estimates in Table 3 and the second-stage estimates that follow control for year of birth, state of birth, and month of birth.

order of 150 for whites and 10 for nonwhites.[17]

3.3.3 Mortality and Survivor Bias in the Census Sample

As a preliminary step, we looked for under-representation of draft-eligible men in the census sample. This analysis is motivated by the possibility that draft-eligible men were more likely to have been killed in wartime and by the link between Vietnam-era service and civilian mortality established in the Hearst, Newman, and Hulley (1986) draft-lottery study of the long-term consequences of Vietnam-era service. Following the mortality investigation, which shows little evidence of an impact of Vietnam-era service on mortality, we look at the effects on self-reported disability rates and labor force status.

Mortality effects are of interest both as an important health outcome and because excess mortality among draft-eligible men may induce selection bias in samples of survivors. The two most likely channels for excess mortality among draft-eligible men are war-related deaths and elevated post-service mortality. The latter may be due to physical injury, PTSD, or other long-term consequences of military service, such as an increased likelihood of cigarette smoking as found by Bedard and Deschenes (2006) for World War II veterans. The excess deaths in the Hearst, Newman and Hulley (1986) study are due to suicide and motor-vehicle accidents, both of which have been linked to PTSD.

Roughly 47,000 men died as a result of hostile action in the Vietnam Era (1964-75) while 8.7 million personnel served in the military during this period for an overall casualty rate of about half a percent. Overall casualty rates among Vietnam-era veterans were low, in part because less than half of active duty personnel served in Indochina, and because of those who did, many served in positions not exposed to combat. Although casualty rates among draftees were higher than the overall Vietnam era death rate, draftees accounted for a minority of combat deaths. It is also noteworthy that over 80 percent of combat deaths occurred before 1970.[18] It therefore seems unlikely that war-related deaths have a large effect on the composition of the sample used in our study.

As a simple check on the possibility of mortality-related selection bias—a potential

[17] A larger instrument set with dummies for RSN 1-30 and RSN 31-60 adds little to the precision obtained with $5zx$. See Angrist and Chen (2008) for more on the draft lottery first stage.

[18] Service and casualty statistics in this paragraph are from Table 583 in the 2000 Statistical Abstract, available on-line at http://www.census.gov/prod/2001pubs/statab/sec11.pdf. Data on casualties by year are available from the national archives: http://www.archives.gov/research/vietnam-war/casualty-statistics.html#year. Statistics on service in Indochina and exposure to combat are from Hearst, Newman, and Hulley (1986).

threat to the validity of using draft eligibility as an instrument—we compared the actual and expected numbers of draft-eligible men in the 2000 Census by race and year of birth. The expected ratio was computed using monthly birth totals for males by race (Vital Statistics Division, 1948-1955), assuming birthdays (and hence lottery numbers) are uniformly distributed within a month. On the whole, draft-eligible men are represented in the census sample almost exactly as predicted, assuming a uniform distribution of lottery numbers within a month. Among whites, the predicted proportion eligible is .407, while the empirical proportion eligible is .405. Among nonwhites, the empirical proportion eligible is slightly more than predicted, .408 versus .405.

Comparisons by single year of birth for white men born 1948-53, reported in appendix Table A1, show draft-eligible men slightly over-represented in three cohorts and slightly under-represented in the other three cohorts (one of these is the 1953 cohort, with no draftees). Only two cohort-specific differences for whites are significant, and all are small. Two out of six cohort-specific contrasts are significant for nonwhites, with slightly more eligibles in the sample than predicted for nonwhites born in 1950 and 1952. Given the magnitudes and sign pattern in this set of comparisons, it seems unlikely that differential mortality by draft-eligibility status had a substantial effect on the composition of the 2000 Census sample. These results also weigh against the view that Vietnam-era service led to elevated civilian mortality.[19]

3.4 Results

3.4.1 Effects on Disability, Transfer Income, and Work

Our main focus is on the effects of Vietnam-era service on self-reported disability status, disability-related transfers, and labor force status, all denoted by Y_i. The empirical framework for these estimates is the equation:

$$Y_i = X_i'\gamma_0 + \beta_0 VET_i + \varepsilon_i, \tag{3.1}$$

[19]The appendix table also breaks out draft-eligibility rates by high school graduation status. The estimates by schooling group also hover around the theoretical proportions, and the rates for HS dropouts are very close to the rates for HS graduates, with significant differences in only two out of six cohorts for whites, and no significant differences in any cohorts for nonwhites. It seems fair to say there is little consistent evidence that draft-eligible men in either schooling group are especially likely to be missing.

where X_i is a vector of controls for state, year and month of birth, and VET_i indicates Vietnam-era veteran status. We construct OLS and 2SLS estimates of (3.1), the latter using the first stages reported in Table 3. As noted in the introduction, veterans may have suffered long-term combat injuries, either physically or as a result of PTSD. Many Vietnam veterans have also been concerned about exposure to the Agent Orange defoliant used by American forces. The loss of earnings associated with Vietnam-era conscription for white veterans (documented in Angrist, 1990) may have also been debilitating, although Angrist and Chen (2008) showed that by 2000 the effect on earnings had diminished. Any lasting health effects should turn up in positive effects on disability for veterans, as captured by census self-reports. In addition, veterans may be more likely to describe themselves as disabled as a consequence of qualification for VDC and/or SSDI. This sort of endogenous disability reporting is discussed by Bound and Waidmann (1992), and Benitez-Silva, Buchinsky, Chan, Cheidvasser, and Rust (2000), among others. Finally, poor health and transfer income may directly affect employment, though each for different reasons, as we discuss further below.

The 2SLS estimates of effects on disability outcomes in panel A of Table 4 suggest that Vietnam-era conscription induced a small overall increase in self-reported disability rates among whites. The estimated effects, reported in columns (3)-(4), range from .012-.014 and are only marginally significant. These effects come from an increase in non-work-limiting disabilities; the estimated effects on overall work-limiting disability rates are nearly zero. There is little evidence of an increase in disability rates for nonwhites, though the 2SLS estimates for nonwhite men, mostly negative, are imprecise. It's also worth noting that the OLS estimates show increased disability rates for whites, while those for nonwhites show a decrease. While the sign pattern of the OLS estimates is consistent with that of the 2SLS estimates, OLS estimates are especially hard to interpret in this context, since men with disabilities are typically precluded from military service.

In contrast with the modest estimated impacts on overall disability rates, panel B shows a marked increase in the likelihood that (both white and nonwhite) veterans receive other income (mostly VDC). The 2SLS estimates here are around .04, which can be compared with a mean proportion receiving other income of .06-.08. At the same time, we find little evidence that Vietnam-era military service raised the proportion of men receiving income from Social Security (mostly SSDI, as noted before). The estimated effect on Social Security income receipt is close to zero for whites, and negative and not significantly different from zero for nonwhites. The estimated effects on the probability of receiving any federal transfers are generally similar to the effects on the probability of receiving other income, reflecting the absence of an overall effect on Social Security income and SSI (results for the latter are not shown in the table).

The largest and most consistent result coming out of an analysis of overall effects on specific disabilities is an increase in the likelihood of a vision or hearing-related problem. These effects, reported in panel C of Table 4, range from .011 for whites to .039 for nonwhites, and both are statistically significant at conventional levels. This increase might reflect an increased incidence of hearing loss or tinnitus among veterans. In contrast, Vietnam-era service does not appear to have worsened average mental health as reflected in the rate at which men report difficulties learning, remembering, or concentrating. This is surprising given the large numbers of Vietnam veterans receiving VDC for PTSD and the fact that such difficulties are recognized PTSD symptoms (Institute of Medicine and National Research Council, 2007).

The veteran effect on non-work-limiting disability rates for whites may reflect a negative causal impact of military service on health. However, the pattern of disability effects does not seem consistent with an interpretation of the increase in disability transfers as the downstream consequence of poor veteran health. First, even if we ignore the work/non-work distinction, the effects of military service on disability rates are too small to explain the increase in disability-related transfers. In other words, if military service affected disability-related transfer receipt only through service-induced disabilities, one could consider military service as an implicit instrumental variable for the effect of disability status on transfer receipt. However, since the transfer effect is around 4 percentage points while the disability effect is a little over one percentage point, this would imply impossibly large (i.e., greater than unity) estimates of the effect of disability status on transfer receipt. Military service must therefore affect transfer receipt for reasons other than disability status (such as financial incentives). In fact, the effects on work-limiting disability—which would seem the most likely to have health consequences that qualify veterans for disability transfers—are nearly zero. Finally, the disability effects in Table 4 do not appear to have translated into lower employment rates or reduced labor force participation, as would usually be expected for workers with consequential health limitations.

The next section further explores the link between Vietnam-era military service and disability, focusing on how this link varies with earnings potential.

3.4.2 Interactions with Predicted Wages and Schooling

The empirical literature on the unintended economic consequences of disability insurance has two themes. The first is that such programs increase the likelihood of early retirement. For example, Bound and Waidmann (1992) and Stapleton and Burkhauser (2003) present evidence suggesting that disability insurance contributed significantly to the drop in labor force participation of near-elderly men over the second half of the twentieth century. A second strand of this literature argues that disability insurance has become increasingly attractive for (non-elderly) low-skilled men because declining real wages for the less skilled have meant a rise in disability insurance replacement rates. In particular, Autor and Duggan (2003) find a close link between enrollment in SSDI or SSI and regional variation in wage levels. Black, Daniel, and Sanders (2002) similarly show that disability insurance take-up rates are highly sensitive to regional variation in labor demand.

As with Social Security disability programs, VDC may provide an attractive alternative to employment for low-wage men even if their disabling conditions are not serious enough to prevent or limit paid employment. In support of this view, Duggan, Rosenheck, and Singleton (2006) show that enrollment in the VDC program seems highly sensitive to small changes in eligibility criteria and in unemployment rates. Moreover, paralleling the incentives created by SSDI and SSI, VDC should reduce work for low-wage men through both income and substitution effects. Substitution effects arise because many veterans are awarded benefits at the 100 percent level on the basis of an IU determination that depends in part on low earnings.

We explore the link between earnings potential and disability outcomes for Vietnam veterans by looking for variation in the causal effects of veteran-status across skill groups. If causal effects on VDC take-up rates and self-reported disability status are driven primarily by deteriorating health, we should not expect these effects to be larger for men in the lowest skill groups, unless low-skilled men were also more likely to have suffered wartime injuries. On the other hand, if VDC is used primarily as an alternative to work for those with low earnings potential, we should see a strong gradient in the effects of veteran status.

Our interacted models use predicted wages and a schooling variable to define skill groups. The predicted wage is the fitted value from a regression of non-veterans' weekly wages on state of birth and education interactions, controlling for year of birth, run separately for whites and nonwhites.[20] Descriptive statistics for subsamples classified by predicted wage appear in columns 4-7 of Table 2. Not surprisingly, these statistics show that men with a lower predicted wage are much more likely to be disabled and have reduced labor force attachment. As a robustness check, we report results from an alternative specification using interactions classified by four schooling groups only. This generates a scheme that can be matched to our analysis of combat exposure by schooling group, described below. As it turns out, the two classification schemes for interacted models produce similar results.

The empirical framework for models with interactions is

$$Y_i = \sum_{j=1}^{4} (\alpha_j D_{ij} + \beta_j D_{ij} VET_i) + X_i'\gamma + \varepsilon_i, \tag{3.2}$$

where the variables D_{i1} to D_{i4} are indicators either for the four schooling groups, or for men with a predicted wage below the 10th percentile, between the 10th and 25th percentile, between the 25th and 75th percentile, and above the 75th percentile. The $D_{ij} \times VET_i$ terms are treated as endogenous and a set of four $D_{ij} \times ELIG_i$ terms are used as excluded instruments in 2SLS estimation of a just-identified model (Table 4 suggests little precision is gained in over-identified models). As before, the vector of covariates, X_i, contains dummies for year, month, and state of birth. The coefficients of interest are β_1 to β_4, the estimated causal effect of Vietnam veteran status on men in each predicted wage or schooling group.[21]

The resulting 2SLS estimates of veteran effects by skill level appear in Table 5. Column 1 in Panel A shows that for white men in the lowest wage group, the effect of Vietnam veteran status on the probability of reporting any disability is .109 (the mean disability rate in this group is .4, as shown in Table 2). There are also much smaller, though still

[20] The sample for the predicted wage regressions consists of male US-born non-veterans born 1948-52. The dependent variable is the weekly wage and the explanatory variables are a full set of state of birth by education-group effects plus year of birth main effects. The education groups are: high school dropout, high school graduate, some college, and college graduate, as described in the appendix. Wage prediction regressions were run separately for whites and nonwhites.

[21] The schooling classification scheme also provides a check on selection bias from the possible endogeneity of schooling. Specifically, 2SLS estimates using the draft lottery (reported in Angrist and Chen, 2008) show that Vietnam-era military service increased the likelihood of college attendance but had little effect on schooling at the high school level or below. Therefore, draft-lottery estimation of veteran effects conditional on high school graduation status is unaffected by any post-treatment selection bias that might contaminate contrasts by college graduation status. This in turn means that the difference in treatment effects between high school dropouts and the other three groups is not subject to bias from conditioning on an outcome variable.

marginally significant, effects of .029 and .018 on white men with wages in the next two quantile groups, but no effect on men with high earnings potential. An almost identical pattern arises when predicted wage groups are replaced by schooling groups, as can be seen in Panel B.

The overall disability effects are decomposed into effects on work-limiting and non-work-limiting disabilities in columns 2 and 3 of Table 5. Consistent with a similar breakdown in the full sample, the first row in each panel of Table 5 shows that the large veteran effect on any disability for low-skilled men is due mostly to an effect on non-work-limiting disabilities, with no significant effect on work-limiting disabilities in any group, though the point estimate for the effect on work-limiting disabilities for men with wages in the lowest group is still substantial. For the highest skill group, the estimated veteran impacts on both non-work and work-limiting disabilities are essentially zero.

Effects on transfer income are generally somewhat larger than those on non-work-limiting disabilities, as shown in columns 4-6. Moreover, while veterans at all predicted wage levels are estimated to be more likely to receive other income (mostly VDC), the largest effect is again for men with the lowest earnings potential. The estimated effects on other income in the low skill groups are .069 using predicted wages and .08 using schooling groups. The effect of veteran status on the likelihood of receiving Social Security income (SSDI) is smaller than that on other income but still significant for men in the lowest skill groups. This suggests that many men leaving the labor force to receive VDC also qualify for and receive SSDI, as argued by Autor and Duggan, 2009).

On the other hand, the results for any federal transfers indicate that SSDI (and to a lesser extent, SSI) is a partial substitute for VDC. This is apparent from the fact that effects on the aggregate transfer category are larger than the effects on any single component. For example, veterans in the lowest skill group about 10 percentage points more likely to receive federal transfers, while those in the next lowest group are about 5 percentage points more likely to receive transfers. Given that VDC and SSDI are two separate programs with independent disability determination procedures, disabled veterans may begin receiving SSDI while they are waiting (or trying) to qualify for VDC or vice versa. Although SSDI is not especially designed to be attractive to veterans, all veterans who apply are required to submit their military discharge papers (form DD-214). In practice, military service increases SSDI benefits and the likelihood of SSDI qualification for men with weak labor force attachment because time in the military generates earnings credits in addition to base pay.[22]

The last set of results in Table 5 is for the veteran effect on employment and labor force participation. These estimates, reported in columns 7 and 8, show a marked decrease in employment and labor force participation among men in the lowest skill groups, with more muted effects in the middle of the predicted wage or schooling distribution, and no effect for men at the top of these distributions. The parallel between the variation in employment effects across skill groups and the pattern of effects on disability and transfer income is striking. However, because the veteran effect on unemployment or being out of the labor force exceeds the effect on work-limiting disability reported in column 2, especially for the least skilled men, disability-induced work limitations seem unlikely to be the sole explanation for reduced veteran employment.[23]

Finally, Table 6 looks at effects on specific disability types. The estimated impact of Vietnam-era service is, again, largest at the low end of the predicted wage or schooling distribution, with no significant effects at the high end. The largest impact at the low end is on physical disabilities, a category that probably includes most muscular and skeletal problems (e.g., related to knees or back). There appear to be smaller effects on physical disabilities in the second-lowest skill group. The second largest set of veteran effects relates to mental disabilities, including difficulties in learning, remembering and concentrating. Vision and hearing problems also appear to have been aggravated by military service among men at the low end of the skill distribution, though the estimates are less precise. Thus, within skill groups, the impact of military service on specific limitations seems broadly consistent with the diagnoses most prevalent among VDC claimants, seen in Table 1.[24]

[22]See, e.g., the pamphlet *Social Security and Military Service,* available at http://www.ssa.gov/pubs/10017.pdf.

[23]Much as we argued in Section 4.1 regarding the relation between effects on disabilities and transfers in the full sample, if the only channel whereby military service affects employment and labor force participation were work-limiting disabilities, the effect of military service on work-limiting disabilities would necessarily be larger than the effect on employment outcomes. The fact that the results come out otherwise suggests something other than health contributes to the employment effects.

[24]We also looked at models for log wages as in Angrist and Chen (2008). The estimated effects on log wages, both overall and within skill groups are essentially zero. We also tried a model that allows for interactions between veteran status and state-level VDC generosity as measured by the fraction of veterans designated IU or with a 100 percent disability rating. This generated marginally significant positive interactions in equations for non-work-limiting disability for both whites and nonwhites and significant interactions in equations for any- and work-limiting disability for nonwhites. On the other hand, we cannot really say whether these interactions reflect the differential application of VA policy or the unobserved characteristics of the veterans who reside in different states.

3.5 Interpreting the Impact of Vietnam-era Service Across Skill Groups

A natural question raised by the results in Table 5 is whether the effects of Vietnam-era service on disability, transfer income, and employment for low-skilled men can be accounted for by differences across skill groups in exposure to combat or the risk of service-related injury. For example, to be diagnosed with PTSD, a veteran must establish that he was exposed to traumatic events of an extreme nature (VA Office of the Inspector General, 2005, p. 46). We therefore ask whether less-educated men were more likely to be exposed to combat or war or to have suffered a service-connected disability. We explore this question using the 1987 Survey of Veterans (known as the SOV-III since it was the third in a series of veteran surveys). The SOV-III interviewed veterans (excluding those still on active duty) in CPS outgoing rotation groups from April 1986 through January 1987. The survey covered roughly two thousand Vietnam veterans and collected information on veterans' service experiences and health. Most relevant for us, the SOV-III included questions about service location and exposure to combat, as well as a direct assessment of service-connected disabilities. The SOV-III asks specifically about service-connected injuries and disabilities, while the disability variables in the 2000 Census are more general. As expected, the disability rates observed in the Census are larger. The data appendix describes the definitions of the variables and the criteria used to select our extract, which is a subset of the sample analyzed in Angrist (1993). Because the results in Tables 4 and 5 show significant effects only for whites, we focus on white men in the SOV-III.

Among all white Vietnam veterans in our SOV-III extract, 40 percent report having served in the Vietnam War theatre (Vietnam, Laos, or Cambodia), 36 percent report exposure to combat, 46 percent report exposure to combat or war, and 6.1 percent report a service-connected disability.[25] To increase the sample size, we analyze an extract that includes men born 1943-57 in addition to an extract limited to the draft lottery cohorts (men born 1948-52). The descriptive statistics for both samples are broadly similar, as can be seen by comparing the descriptive statistics in the first rows of Panels A and B in Table

[25]Respondents are coded as having been exposed to combat if they responded in the affirmative to a question asking whether they were in or exposed to combat. Respondents are coded as having been exposed to combat or war if they indicated that they were either exposed to combat or were stationed in a war zone. Respondents with a service-connected disability are those who indicated they have been notified by the VA that they are eligible for payment for a service-connected medical condition or disability.

7.

Our empirical analysis of the relationship between education and war exposure is structured by regressions of combat or war exposure and service-connected disabilities on schooling dummies similar to those used to construct the estimates in Panel B of Table 5.[26] Specifically, the schooling dummies are indicators for high school graduates, men with some college, and college graduates, with high school dropouts as the omitted reference group. The estimated coefficients on the schooling dummies, reported in columns (1)-(4) of Table 7, show little difference in the likelihood of combat/war exposure or service-related disability across schooling groups. For example, while 33 to 40 percent of veterans with no high school diploma reported ever being exposed to combat, as shown in Panel A, column (2), the combat exposure rates were only 1 percentage point less for men with some college or a college degree.

Because the largest effects of Vietnam veteran status on disability and employment (reported in Table 5) appear among high school dropouts, the relationship between high school dropout status and exposure to war or combat is of special interest. Estimates of this relationship in columns 5-7 of Table 7 show a small and insignificant relationship between having a high school diploma and the likelihood of serving in the war theatre, or being exposed to combat or war.

The muted relationship between education and exposure to war or combat does not support the notion that less educated veterans were more likely to have been exposed to war-related trauma or injury. At the same time, we note the possibility, investigated by Macklin, Metzger, Litz, McNally, Lasko, Orr, and Pitman (1998) (among others), that exposure to the same traumatic experience may be more likely to trigger PTSD in veterans with lower cognitive ability (or less education). Moreover, exposure to combat or service in a war zone may have entailed different experiences for veterans of different education levels (because of, say, differences in rank). We therefore look at a direct measure of service-connected disability and examine how this measure varies with education. Columns 4 and 8 in Table 7 show no significant difference in the likelihood of a service-related disability across schooling groups, either in the full sample of Vietnam veterans (Panel A), or in the draft lottery subsample (Panel B). These findings are inconsistent with the notion that

[26]Definitions of SOV-III schooling groups appear in the data appendix. Schooling group dummies describe education at the time of the SOV-III survey. The covariates in these regressions consist of 5-year cohort dummies.

less-educated men were especially vulnerable to PTSD or other service-connected injuries.

Although the results in Table 7 are somewhat imprecise and therefore less than conclusive, they are consistent with the findings in Table 5 in pointing away from health *per se* as the primary explanation for lower employment rates among Vietnam veterans with low levels of education or low earnings potential. As noted earlier, the effects of Vietnam-era service on work-limiting disabilities are too small (on the order of 4-5 percentage points) to account for employment reductions ranging from 8-12 percentage points (depending on the outcome and skill group definition). It therefore seems likely that part of the explanation for service-induced increases in disability rates among the low-skilled is an *ex post* validation of VDC or SSDI eligibility, a status administratively bound up with employment. Specifically, a veteran who qualifies for a federal disability insurance program may be more likely to identify himself as disabled, even if his disability does not limit work. The strong effect of Vietnam-era veteran status on aggregate transfer income, reported in column 6 of Table 5, seems to be the leading proximate cause for the negative effect of veteran status on employment and labor force participation rates among low-skilled men.

3.6 Conclusions

Our estimates of the causal effects of Vietnam-era military service on disability rates, transfer income, and employment paint a complicated picture. We find only a small service-induced increase in overall disability rates among white veterans, and an insignificant decrease among nonwhites. Moreover, the increase among whites comes almost entirely from disabilities judged by census respondents not to be work-limiting. At the same time, an analysis of effects by skill groups, using either predicted wages or schooling, shows a sizeable effect on disability among the least skilled white veterans, with some smaller but still significant effects in the lower-middle of the skill distribution. We also find large negative effects on employment in the lowest skill groups.

Did the least skilled suffer the most serious and lasting health consequences of Vietnam-era service? Our analysis points away from this interpretation. First, less-educated men were not more likely to serve in the Vietnam War theatre, to be exposed to combat or war, or to have reported a service-connected disability in 1987. In addition, the estimated effects of Vietnam-era veteran status on work-limiting disabilities are too small to explain the estimated effects on employment and labor force participation. A case can therefore be made for disability insurance as a primary causal agent driving these results, even allowing for a modest negative overall health effect suggested by our estimates. Veterans who get VDC (or SSDI), especially those who are (or aspire to be) classified as "individually unemployable," are probably more likely to define themselves as disabled and less likely to work. This seems to be a special concern for Vietnam-era PTSD claims; data from 2005 show that roughly one-third of PTSD claimants are designated IU and that IU claimants are concentrated in the Vietnam cohort (Christensen, et al, 2007, Figures 58-59).

Our results have important implications for veterans compensation policy. The number of Vietnam-era VDC beneficiaries grew rapidly in the late 1990s, growth that accelerated in the early part of this century and has not yet leveled off. This imposes a growing burden on a system that must serve new cohorts of veterans from the Gulf War, Afghanistan, and Iraq. The results reported here suggest the growth in Vietnam-era disability claims (and hence costs) are not only a manifestation of the health consequences of the Vietnam war, but also a reflection of the incentives embedded in our disability insurance system for veterans. While our estimates are specific to those individuals whose veteran status was determined by their draft eligibility status, the incentives in the compensation programs likely apply to veterans more broadly.

Our findings also raise questions about widely publicized projections of the disability costs likely to come out of current conflicts. Specifically, Stiglitz and Bilmes (2008, pp. 82-83) note that a large number of VDC claims in this most recent cohort are for PTSD and that PTSD is an especially expensive diagnosis associated with high program costs and large earnings losses. But the costliness of PTSD claims comes in large part from the link with IU and the consequent increase in VDC benefits. Case reviews in VA Office of the Inspector General (2005) show that mental health visits declined by 82 percent after an IU rating decision, and that many granted IU status stop seeking treatment for mental health entirely, though health care visits for other conditions are unchanged. Likewise, our results indicate that the employment consequences of PTSD may have as much to do with incentives as with a medical inability to work, at least in many cases. The complicated links between military service and variables related to health show that the disability-related costs of conflict are driven by policy and regulatory choices, as well as the battlefield consequences of war.

References

ANGRIST, JOSHUA D. (1990): "Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records," *American Economic Review*, 80, 313–335.

——— (1993): "The Effect of Veterans Benefits on Education and Earnings," *Industrial and Labor Relations Review*, 46, 637–652.

ANGRIST, JOSHUA D., AND STACEY H. CHEN (2008): "Long-Term Economic Consequences of Vietnam-Era Conscription: Schooling, Experience and Earnings," IZA Discussion Paper No. 3628.

ANGRIST, JOSHUA D., GUIDO IMBENS, AND DONALD B. RUBIN (1996): "Identification of Causal Effects Using Instrumental Variables," *Journal of the American Statistical Association*, 91, 444–472.

AUTOR, DAVID H., AND MARK G. DUGGAN (2003): "The Rise In The Disability Rolls And The Decline In Unemployment," *Quarterly Journal of Economics*, 118, 157–205.

——— (2008): "The Effect of Transfer Income on Labor Force Participation and Enrollment in Federal Benefits Programs: Evidence from the Veterans Disability Compensation Program," Unpublished manuscript.

BEDARD, KELLY, AND OLIVIER DESCHENES (2006): "The Long-Term Impact of Military Service on Health: Evidence fron World War II and Korean War Veterans," *American Economic Review*, 96.

BENITEZ-SILVA, HUGO, MOSHE BUCHINSKY, HIU MAN CHAN, SOFIA CHEIDVASSER, AND JOHN RUST (2000): "How Large is the Bias in Self-Reported Disability?," NBER Working Paper No. 7526.

BLACK, DAN, KERMIT DANIEL, AND SETH SANDERS (2002): "The Impact of Economic Conditions on Participation in Disability Programs: Evidence from the Coal Boom and Bust," *The American Economic Review*, 92, 27–50.

BOUND, JOHN, AND TIMOTHY WAIDMANN (1992): "Disability Transfers, Self-Reported Health, and the Labor Force Attachment of Older Men: Evidence from the Historical Record," *The Quarterly Journal of Economics*, 107, 1393–1419.

CHRISTENSEN, ERIC, JOYCE MCMAHON, ELIZABETH SCHAEFER, TED JADITZ, AND DAN HARRIS (2007): "Final Report for the Veterans' Disability Benefits Commission: Compensation, Survey Results, and Selected Topics," Alexandria, Virginia: CNA Corporation.

COOPER, DANIEL L. (2005): "Statement Before the Senate Committee on Veterans' Affairs, October 27, 2005," Available at http://www.va.gov/OCA/testimony/svac/05102720.asp.

DOBKIN, C., AND R. SHABANI (2006): "The Long Term health Effects of Military Service: Evidence from the National Health Interview Survey and the Vietnam Era Draft Lottery," University of California at Santa Cruz, Department of Economics, mimeo.

DUGGAN, MARK, ROBERT ROSENHECK, AND PERRY SINGLETON (2006): "Federal Policy and the Rise in Disability Enrollment: Evidence for the VA's Disability Compensation Program," NBER Working Paper 12323.

DUGGAN, MARK, PERRY SINGLETON, AND JAE SONG (2007): "Aching to Retire? The Rise in the Full Retirement Age and its Impact on the Social Security Disability Rolls," *Journal of Public Economics*, 91, 1327–1350.

EISENBERG, DANIEL, AND BRIAN ROWE (2008): "The Effects of Smoking in Young Adulthood on Smoking and Health Later in Life: Evidence Based on the Vietnam Era Draft Lottery," US Census Bureau Center for Economic Studies Paper No. CES-WP-08-35.

FRUEH, B. CHRISTOPHER, JON D. ELHAI, ANOUK L. GRUBAUGH, JEANNINE MONNIER, TODD B. KASHDAN, JULIE A. SAUVAGEOT, MARK B. HAMNER, B.G. BURKETT, AND GEORGE W. ARANA (2005): "Documented combat exposure of US veterans," *British Journal of Psychiatry*, 186, 467–472.

GOLDBERG, J., M. RICHARDS, R. ANDERSON, AND M. RODIN (1991): "Alcohol Consumption in Men Exposed to the Military Draft Lottery: A Natural Experiment," *Journal of Substance Abuse*, 3, 307–313.

HEARST, N., T. BUEHLER, T. NEWMAN, AND G. RUTHERFORD (1991): "The Draft Lottery and AIDS: Evidence Against Increased Intravenous Drug Use by Vietnam Veterans," *American Journal of Epidemiology*, 134, 522–525.

HEARST, N, TB NEWMAN, AND SB HULLEY (1986): "Delayed Effects of the Military Draft on Mortality. A Randomized Natural Experiment," *New England Journal of Medicine*, 314, 620–624.

INSTITUTE OF MEDICINE AND NATIONAL RESEARCH COUNCIL (2007): *PTSD Compensation and Military Service*. The National Academies Press, Washington, DC.

IVERSEN, AMY, TRUDIE CHALDER, AND SIMON WESSELY (2007): "Gulf War Illness: Lessons from Medically Unexplained Symptoms," *Clinical Psychology Review*, 27, 842 – 854.

MACKLIN, MICHAEL L., LINDA J. METZGER, BRETT T. LITZ, RICHARD J. MCNALLY, NATASHA B. LASKO, SCOTT P. ORR, AND ROGER K. PITMAN (1998): "Lower precombat intelligence is a risk factor for posttraumatic stress disorder," *Journal of Consulting and Clinical Psychology*, 66, 323–326.

MEDICAL RESEARCH COUNCIL (2003): "MRC Review of Research into UK Gulf Veterans' Illnesses," London: The Medical Research Council.

PHILPOTT, TOM (2005): "Too many 'IU' vets?," *Stars and Stripes*, Pacific edition, Thursday, May 19, 2005. Accessed from http://www.stripes.com/article.asp?section=104&article=29174 on December 15, 2008.

RESEARCH ADVISORY COMMITTE ON GULF WAR VETERANS' ILLNESSES (2008): "Gulf War Illness and the Health of Gulf War Veterans: Scientific Findings and Recommendations," Washington, D.C.: U.S. Government Printing Office.

Rosenheck, Robert A., and Alan F. Fontana (2007): "Recent Trends In VA Treatment Of Post-Traumatic Stress Disorder And Other Mental Disorders," *Health Affairs*, 26, 1720–1727.

Seltzer, Carl C., and Seymour Jablon (1974): "Effects of Selection on Mortality," *American Journal of Epidemiology*, 100.

Social Security Administration (2009): "Social Security and Military Service," SSA Publication No. 05-10017, Available at http://www.ssa.gov/pubs/10017.pdf.

Stapleton, David C., and Richard V. Burkhauser (eds.) (2003): *The Decline in Employment of People with Disabilities: A Policy Puzzle.* W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Stiglitz, Joseph, and Linda Bilmes (2008): *The Three Trillion Dollar War.* Allen Lane.

U.S. Bureau of the Census (2000): "Statistical Abstract of the United States, 2000," Washington, DC: U.S. Government Printing Office.

——— (2006): "Statistical Abstract of the United States, 2006," Washington, DC: U.S. Government Printing Office.

U.S. General Accountability Office (2006): "Veterans' Disability Benefits: VA Should Improve Its Management of Individual Unemployability Benefits by Strengthening Criteria, Guidance, and Procedures," GAO Report GAO-06-309.

U.S. General Accounting Office (1987): "Improving the Integrity of VA's Unemployability Compensation Program," General Accounting Office Report HRD-87-62.

VA Office of the Inspector General (2005): "Review of State Variances in VA Disability Compensation Payments," VA Office of Inspector General Report No. 05-00765-137, Washington, D.C.

Veterans Benefits Administration (2007): "Annual Benefits Report for Fiscal Year 2006," Washington, D.C.: Veterans Benefits Administration.

Vital Statistics Division, Bureau of the Census (1948-1955): "Vital Statistics of the United States," Washington, D.C.: United States Government Printing Office.

3.7 Tables and Figures

Table 1: Veterans disability compensation claims: Most common diagnoses for Vietnam veterans in 2005

Year	Veteran Population (1)	Disabilities per Veteran (2)	Veterans receiving compensation for: Any disabilities (3)	Diabetes* (4)	PTSD (5)	Hearing (loss of acuity) (6)	Scars (7)	Generalized Musculo-skeletal Conditions (8)	Hyper-tension (9)	Arthritis (due to trauma) (10)	Knee impairment (11)
						A. Vietnam					
1999	8,113,000	2.76	735,627	-	90,695	50,184	127,023	82,446	56,231	48,263	67,836
2001	7,916,774	2.77	749,554	-	106,809	60,753	125,939	80,586	55,545	53,332	66,335
2003	8,210,925	2.88	848,156	135,011	142,876	95,931	125,534	78,413	66,084	62,821	65,026
2005	8,054,993	3.00	916,220	190,199	179,737	129,323	121,850	78,270	72,169	69,034	62,713
						B. Korea					
1999	4,064,000	2.01	174,807	-	N/A	N/A	18,879	N/A	N/A	8,903	6,862
2001	3,347,310	2.04	166,362	-	6,524	N/A	17,703	6,452	N/A	9,230	6,411
2003	3,580,249	2.12	164,482	-	8,994	15,659	16,761	N/A	N/A	9,941	N/A
2005	3,256,925	2.2	161,512	-	10,994	25,529	15,476	5,552	N/A	10,030	N/A
						C. Gulf War					
1999	2,223,000	3.20	282,140	-	N/A	28,208	33,719	59,337	28,405	35,304	56,320
2001	3,095,952	3.32	365,780	-	N/A	36,399	42,523	77,849	37,260	52,826	63,966
2003	3,783,414	3.48	476,026	-	N/A	47,031	52,479	100,925	49,948	77,202	72,300
2005	4,377,845	3.70	611,729	-	N/A	60,023	60,350	131,092	64,558	100,374	81,677

Notes: This table reports the number of veterans receiving VDC in total and for specific disabilities. The listed diagnoses are the top 10 conditions in 2005, except for tinnitus (which is often diagnosed with loss of hearing acuity) and degenerative arthritis (which is not in the top 10 before 2005). Diabetes recognized as a service-related impairment in 2002. "N/A" denotes that the diagnosis was not among the top ten compensated diagnoses for the given year and service era.

Source: Veterans Benefits Administration Annual Reports for 1999, 2001, 2003, and 2005.
* Diabetes not presumed service-related for Korea- and Gulf War-era service.

Table 2. Descriptive statistics by race and veteran status

	All	Vietnam veteran	Non-veteran	Wage index percentile			
				<10	10-25	25-75	>75
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		A. Whites					
Draft eligibility (by RSN)	0.437	0.552	0.386	0.425	0.424	0.438	0.446
Veteran status (served in Vietnam Era)	0.305	1.000	0.000	0.261	0.354	0.338	0.228
Age	49.2	49.6	49	49.1	49.1	49.2	49.3
		i. Disability variables					
Any disability	0.198	0.217	0.190	0.386	0.265	0.183	0.113
Work-limiting disability	0.124	0.134	0.120	0.243	0.172	0.113	0.070
Non-work-limiting disability	0.074	0.083	0.070	0.143	0.093	0.070	0.043
		ii. Transfer income					
Other income (mostly VDC) > 0	0.059	0.109	0.037	0.069	0.068	0.062	0.044
SSA income excluding SSI (mostly SSDI) > 0	0.035	0.036	0.034	0.097	0.051	0.028	0.014
SSI > 0	0.017	0.013	0.019	0.060	0.024	0.012	0.005
Any Federal transfer income > 0	0.100	0.142	0.082	0.197	0.129	0.093	0.059
		iii. Specific disability types					
Mental (difficulty learning, remembering, or concentrating)	0.045	0.049	0.044	0.134	0.062	0.036	0.019
Vision or hearing (blindness, deafness, or a severe vision or hearing impairment)	0.038	0.043	0.036	0.076	0.049	0.035	0.021
Physical (limitation to physical activities e.g. walking, climbing stairs, reaching, lifting, or carrying)	0.086	0.101	0.080	0.193	0.118	0.078	0.041
Mobility (difficulty going outside the home alone)	0.052	0.054	0.052	0.133	0.077	0.044	0.023
Self-care (difficulty dressing, bathing, or getting around inside the home)	0.022	0.022	0.022	0.062	0.030	0.018	0.010
		iv. Labor market variables					
Not working	0.145	0.154	0.141	0.327	0.205	0.125	0.074
Not in labor force	0.118	0.126	0.115	0.281	0.170	0.100	0.056
N	1,141,551	353,367	788,184	114,588	171,459	572,311	283,193
		B. Nonwhites					
Draft eligibility (by RSN)	0.440	0.520	0.406	0.427	0.424	0.443	0.449
Veteran status (served in Vietnam Era)	0.293	1.000	0.000	0.148	0.195	0.337	0.326
Age	49.2	49.5	49	49.0	49.0	49.2	49.2
		i. Disability variables					
Any disability	0.332	0.326	0.334	0.447	0.409	0.338	0.226
Work-limiting disability	0.212	0.205	0.215	0.275	0.253	0.217	0.151
Non-work-limiting disability	0.120	0.120	0.120	0.173	0.156	0.121	0.075
		ii. Transfer income					
Other income (mostly VDC) > 0	0.078	0.140	0.052	0.067	0.069	0.083	0.078
SSA income excluding SSI (mostly SSDI) > 0	0.060	0.056	0.062	0.093	0.085	0.060	0.032
SSI > 0	0.044	0.029	0.051	0.086	0.072	0.041	0.017
Any Federal transfer income > 0	0.163	0.198	0.148	0.218	0.201	0.163	0.116
		iii. Specific disability types					
Mental (difficulty learning, remembering, or concentrating)	0.076	0.072	0.077	0.136	0.115	0.071	0.036
Vision or hearing (blindness, deafness, or a severe vision or hearing impairment)	0.048	0.047	0.049	0.070	0.063	0.047	0.032
Physical (limitation to physical activities e.g. walking, climbing stairs, reaching, lifting, or carrying)	0.139	0.145	0.136	0.205	0.178	0.140	0.085
Mobility (difficulty going outside the home alone)	0.122	0.112	0.126	0.173	0.162	0.124	0.072
Self-care (difficulty dressing, bathing, or getting around inside the home)	0.042	0.036	0.044	0.069	0.061	0.040	0.022
		iv. Labor market variables					
Not working	0.338	0.295	0.356	0.508	0.468	0.341	0.183
Not in labor force	0.284	0.247	0.300	0.440	0.400	0.286	0.148
N	154,810	45,344	109,466	16,002	23,148	77,088	38,572

Notes: This table reports descriptive statistics for men born 1948-52 in the 2000 1:6 census file. Statistics use census sampling weights.

Table 3. First-stage estimates by race and year of birth

	Pooled cohorts	By single year of birth				
	1948-52	1948	1949	1950	1951	1952
	(1)	(2)	(3)	(4)	(5)	(6)
		A. Whites				
Draft-eligibility effect	.112***	.058***	.074***	.133***	.138***	.168***
	(.001)	(.001)	(.003)	(.002)	(.002)	(.002)
RSN effects (5zx):						
RSN 1-95	.128***	.065***	.088***	.154***	.155***	.173***
	(.001)	(.003)	(.003)	(.003)	(.003)	(.003)
RSN 96-125	.082***	.060***	.077***	.131***	.128***	.023***
	(.002)	(.005)	(.005)	(.004)	(.004)	(.003)
RSN 126-160	.058***	.054***	.061***	.126***	.050***	.008***
	(.002)	(.004)	(.004)	(.004)	(.004)	(.003)
RSN 161-195	.044***	.044***	.054***	.102***	.024***	-.001
	(.002)	(.004)	(.004)	(.004)	(.003)	(.003)
RSN 196-230	.006***	.004	.006	.013***	-.001	.008**
	(.002)	(.004)	(.004)	(.004)	(.003)	(.003)
F-statistics	2294	111	202	731	861	1028
N	1,141,551	220,891	224,130	223,984	232,348	240,198
		B. Nonwhites				
Draft-eligibility effect	.072***	.031***	.049***	.090***	.096***	.096***
	(.003)	(.007)	(.006)	(.006)	(.006)	(.006)
RSN effects (5zx):						
RSN 1-95	.081***	.039***	.059***	.101***	.101***	.099***
	(.003)	(.009)	(.008)	(.007)	(.007)	(.007)
RSN 96-125	.058***	.027**	.072***	.089***	.090***	.016*
	(.005)	(.013)	(.012)	(.011)	(.011)	(.009)
RSN 126-160	.041***	.027**	.042***	.093***	.034***	.005
	(.005)	(.012)	(.012)	(.011)	(.010)	(.009)
RSN 161-195	.021***	.012	.027**	.066***	-.005	.005
	(.005)	(.012)	(.011)	(.010)	(.009)	(.009)
RSN 196-230	.001	-.004	.018	.008	-.010	-.006
	(.005)	(.012)	(.011)	(.010)	(.009)	(.009)
F-statistics	134	4.98	14.3	48.9	55.1	47.3
N	154,810	28,272	30,321	31,942	31,162	33,113

Notes: This table reports draft-eligibility and RSN-group effects on the probability of veteran status. Draft-eligibility effects and RSN group effects are from separate regressions. Effects in columns (2)-(6) are from separate regressions by year. Robust standard errors are shown in parentheses. All models include a full set of dummies for state of birth and month of birth, and column (1) also includes year of birth dummies. Statistics use census sample weights. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table 4. OLS and 2SLS estimates of effects on disability and labor force status for men born 1948-52

	Whites				Nonwhites			
			2SLS				2SLS	
Dependent variable	Mean (1)	OLS (2)	elig (3)	5zx (4)	Mean (5)	OLS (6)	elig (7)	5zx (8)
		A. Disability variables						
Any disability	.198	.024***	.012	.014*	.332	-.012	-.061	-.063
		(.001)	(.008)	(.007)		(.003)	(.040)	(.036)
Work-limiting disability	.124	.013***	.000	-.001	.212	-.010	-.045	-.054
		(.001)	(.007)	(.006)		(.003)	(.034)	(.031)
Non-work-limiting disability	.074	.011***	.013**	.014***	.120	-.001	-.016	-.006
		(.001)	(.005)	(.005)		(.002)	(.028)	(.026)
		B. Transfer income						
Other income (mostly VDC) > 0	.059	.072***	.042***	.040***	.078	.087***	.034	.040**
		(.001)	(.005)	(.004)		(.002)	(.022)	(.020)
SSA income excluding SSI (mostly SSDI) >0	.035	.000	.001	.004	.060	-.007	-.027	-.030
		(.000)	(.004)	(.003)		(.002)	(.020)	(.018)
Any Federal transfer income > 0	.100	.058***	.039***	.040***	.163	.047***	.032	.027
		(.001)	(.006)	(.005)		(.002)	(.031)	(.028)
		C. Specific disability types						
Mental	.045	.003***	.007	.006	.076	-.007	.015	.011
		(.001)	(.004)	(.004)		(.002)	(.023)	(.021)
Vision or hearing	.038	.005***	.011***	.012***	.048	-.003	.039**	.036**
		(.000)	(.004)	(.003)		(.001)	(.018)	(.016)
Physical	.086	.018***	.009	.012**	.139	.005**	-.028	-.030
		(.001)	(.006)	(.005)		(.002)	(.029)	(.026)
Mobility	.052	.001**	.005	.005	.122	-.014	-.008	-.008
		(.001)	(.005)	(.004)		(.002)	(.028)	(.025)
Self-care	.022	.000	.007**	.008***	.042	-.009	.011	-.001
		(.000)	(.003)	(.003)		(.001)	(.017)	(.016)
		D. Labor force status						
Not working	.145	.010***	.005	.003	.338	-.063	-.001	-.020
		(.001)	(.007)	(.007)		(.003)	(.040)	(.037)
Not in labor force	.118	.007***	.002	.002	.284	-.057	.026	.016
		(.001)	(.007)	(.006)		(.003)	(.039)	(.035)
N		1,141,551				154,810		

Note: This table reports OLS and 2SLS estimates of the effects of Vietnam veteran status on the dependent variable listed at left. All regressions include a full set of dummies for state of birth, year of birth and month of birth. The estimates in columns 3 and 7 use a simple draft-eligibility dummy as instruments. The estimates in columns 4 and 8 use 5 RSN dummies interacted with year of birth. Robust standard errors are reported in parentheses. Estimates use census sampling weights. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table 5. 2SLS estimates of veteran effects by predicted wage and schooling: Whites

		Disability variables			Transfer income			Labor Force Status	
		Any disability	Work-limiting disability	Non-work-limiting disability	Other income (mostly VDC) >0	SSA income (mostly SSDI) >0	Any Federal transfer income >0	Not working	Not in labor force
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
					A. By wage index percentile				
Veteran status x wage index percentile	<10	.109***	.046	.062**	.069***	.055***	.102***	.081**	.058*
		(.034)	(.030)	(.025)	(.017)	(.021)	(.028)	(.033)	(.032)
	10-25	.029*	.021	.009	.037***	.013	.048***	.038**	.030**
		(.016)	(.014)	(.010)	(.009)	(.008)	(.012)	(.015)	(.014)
	25-75	.018**	.000	.018***	.045***	-.001	.041***	.009	.008
		(.009)	(.007)	(.006)	(.006)	(.004)	(.007)	(.008)	(.007)
	>75	.006	.000	.006	.035***	-.005	.032***	.008	-.008
		(.011)	(.010)	(.008)	(.008)	(.604)	(.009)	(.010)	(.009)
					B. By schooling group				
Veteran status x schooling group	HS dropout	.105**	.039	.066*	.080***	.084***	.134***	.116**	.086*
		(.050)	(.044)	(.037)	(.025)	(.032)	(.042)	(.049)	(.047)
	HS graduate	.036***	.018*	.018**	.033***	.009	.036***	.032***	.027***
		(.012)	(.010)	(.008)	(.007)	(.006)	(.009)	(.011)	(.010)
	Some college	.026**	.009	.017**	.056***	.002	.056***	.014	.009
		(.011)	(.009)	(.008)	(.007)	(.005)	(.009)	(.010)	(.009)
	College degree	.006	-.005	.011*	.036***	-.005	.032***	-.006	-.005
		(.010)	(.008)	(.006)	(.006)	(.004)	(.007)	(.008)	(.007)

Notes: Panel A reports coefficients from a regression in the sample of white men born 1948-1952 of the variable indicated in the column heading on dummies for the wage index percentile and their interactions with Vietnam veteran status, and Panel B reports coefficients from a regression of the variable indicated in the column heading on dummies for education level and their interactions with Vietnam veteran status. The sample size is 1,141,551. All regressions control for state, year, and month of birth. The wage index was computed from a regression of white non-veterans' weekly wages on state of birth and education interactions, controlling for year of birth. Robust standard errors are shown in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table 6. 2SLS estimates of veteran effect interactions for specific disability types: Whites

		Disability Type:				
		Mental (1)	Vision or hearing (2)	Physical (3)	Mobility (4)	Self-care (5)
				A. By wage index percentile		
Veteran status x wage index percentile	<10	.064***	.039**	.088***	.062***	.027
		(.024)	(.019)	(.028)	(.024)	(.017)
	10-25	.014	.022***	.028**	.005	.013**
		(.009)	(.008)	(.012)	(.010)	(.006)
	25-75	.006	.011***	.010*	.008	.010***
		(.004)	(.004)	(.006)	(.005)	(.003)
	>75	.003	.003	-.008	.000	.002
		(.005)	(.005)	(.007)	(.006)	(.004)
				B. By schooling group		
Veteran status x schooling group	HS dropout	.060	.047*	.120***	.056	.036
		(.037)	(.028)	(.042)	(.037)	(.027)
	HS grad	.013**	.021***	.026***	.010	.011**
		(.006)	(.006)	(.009)	(.007)	(.005)
	Some college	.012**	.010*	.015*	.011*	.013***
		(.006)	(.006)	(.008)	(.006)	(.004)
	College degree	.003	.006	-.006	.002	.002
		(.004)	(.004)	(.006)	(.005)	(.003)

Notes: The same as table 5. The sample size is 1,141,551. ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Table 7. Combat and war-theater exposure by education for white Vietnam veterans

	Schooling				HS graduate status only			
	Vietnam/ Laos/ Cambodia	Combat	Combat or war	Service-related disability	Vietnam/ Laos/ Cambodia	Combat	Combat or war	Service-related disability
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				A. Birth cohorts 1943-1957 (N=1893)				
Dependent variable mean	0.398	0.360	0.462	0.089	0.398	0.360	0.462	0.089
Dep. var. mean for HS dropouts	0.406	0.328	0.439	0.112	0.406	0.328	0.439	0.112
Regression estimates								
HS graduate	.021	.050	.054	-.046				
	(.063)	(.063)	(.065)	(.041)				
Some college	-.059	-.010	-.019	-.004				
	(.062)	(.063)	(.066)	(.042)				
College graduate	-.081	-.010	-.044	-.025				
	(.065)	(.066)	(.068)	(.042)				
HS graduate or more					-.033	.014	.004	-.026
					(.059)	(.060)	(.062)	(.040)
F test p-value for education vars.	0.039	0.294	0.071	0.185	0.584	0.818	0.950	0.621
				B. Birth cohorts 1948-1952 (N=724)				
Dependent variable mean	0.417	0.371	0.439	0.093	0.417	0.371	0.439	0.093
Dep. var. mean for HS dropouts	0.404	0.404	0.419	0.095	0.404	0.404	0.419	0.095
Regression estimates								
HS graduate	.072	-.018	.081	-.003				
	(.108)	(.107)	(.108)	(.064)				
Some college	.025	-.007	.029	.001				
	(.108)	(.108)	(.108)	(.063)				
College graduate	-.112	-.113	-.100	-.003				
	(.115)	(.116)	(.116)	(.066)				
HS graduate or more					.014	-.035	.022	-.002
					(.103)	(.103)	(.104)	(.060)
F test p-value for education vars.	0.078	0.461	0.088	0.004	0.893	0.736	0.834	0.000

Notes: the table reports OLS coefficients of the dependent variables in the column headings on dummies for education levels and 5-year cohort dummies in Panel A. The sample in Panel B contains only one 5-year cohort group. The omitted education level is HS dropout. Data are from the third Survey of Veterans (SOV-III), conducted in 1987.

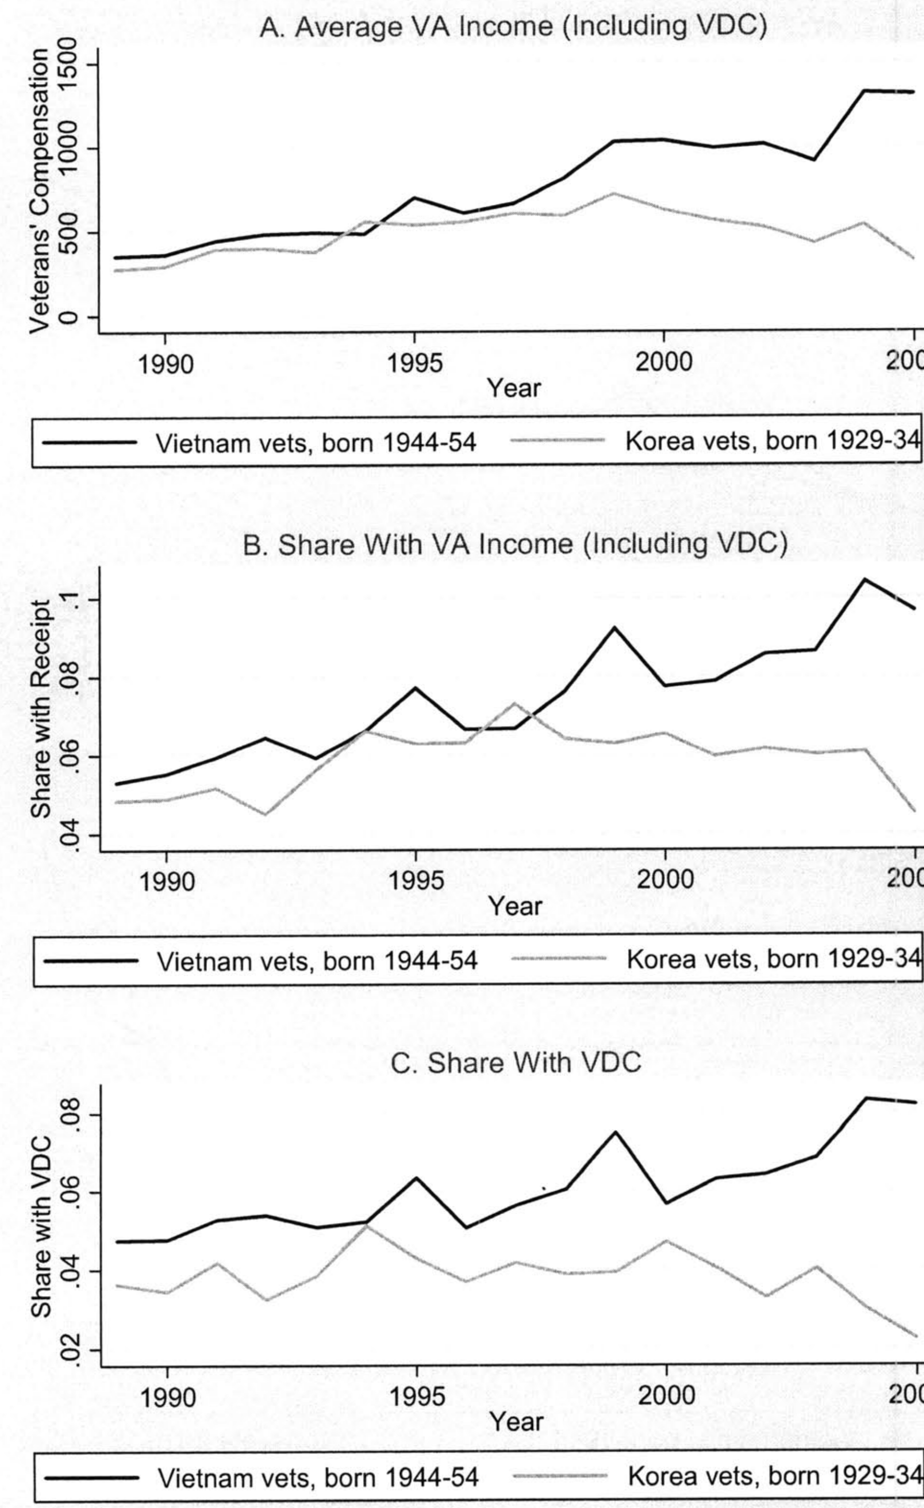

Figure 1. VA Income and Receipt by Year and Service Era

Note: Amounts are in 2005 Dollars. Data are from the March CPS.

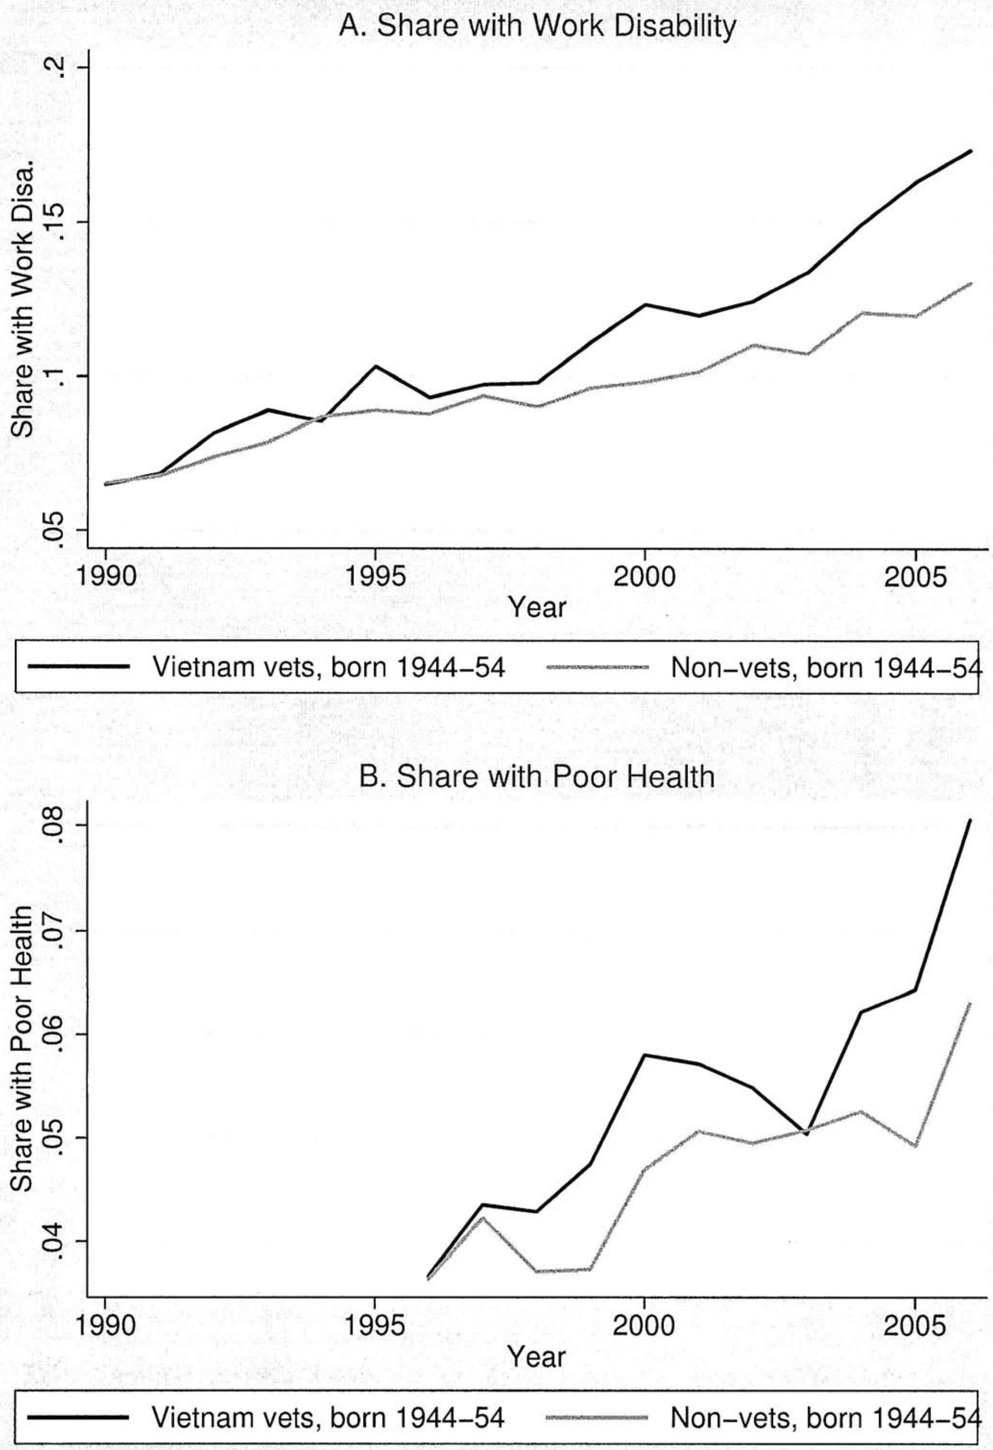

Figure 2. Work Disability and Health by Year and Veteran Status

Note: Data are from the March CPS.

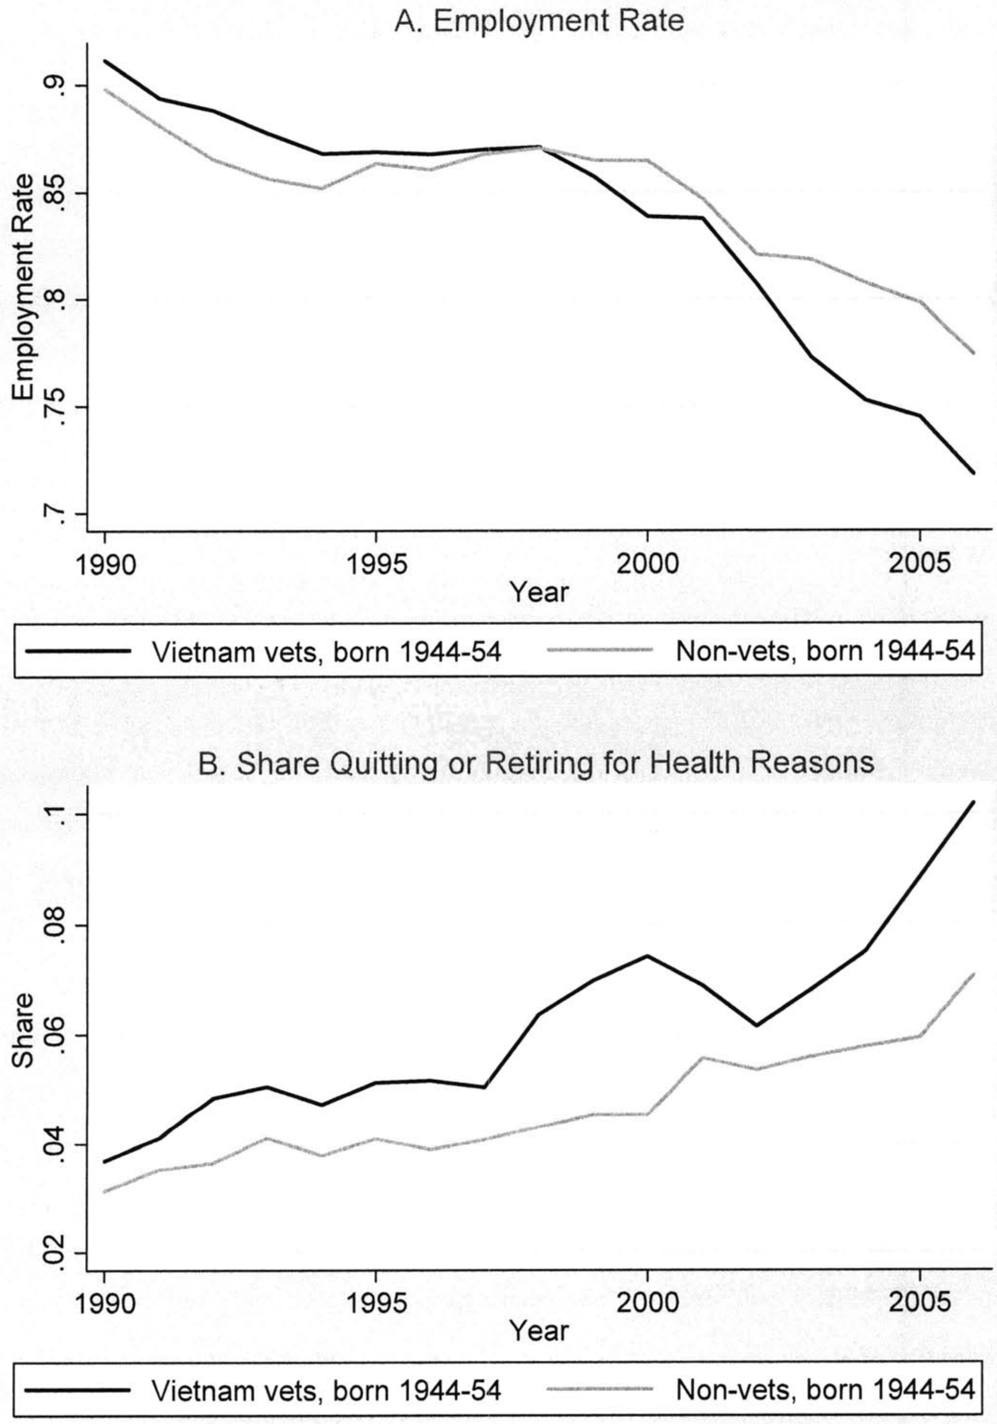

Figure 3. Employment and Health-Related Labor Force Exit by Year and Veteran Status

Note: Data are from the March CPS.

Appendix

Health and Disability Data in the CPS

Figures 1-3 were constructed using data from the 1990-2006 CPS March Demographic Supplements. All data were downloaded from the Minnesota Population Center' s Integrated Public Use Microdata Series (IPUMS), accessible at www.ipums.org. Year of birth was imputed assuming men were born after the survey date. We categorized Vietnam veterans as all men born between 1944-54 who served during the Vietnam era, as reported in the variable VETLAST, which reports an individual' s most recent period of service. Active duty servicemen were excluded. Korean-era veterans were identified the same way, except we used the 1929-1934 birth cohort. non-veterans were classified based on the variable VETSTAT.

The disability-related income variables used in Figure 1 uses the variable INCVET and a dummy for men who received veterans' disability compensation (GOTVDISA), both in 2005 dollars.. INCVET captures any income from the VA, including service-related disability payments (VDC), non-disability pension payments, and educational allowances. GOTVDISA information is collected only for respondents who received veterans' payments during the previous calendar year, and it indicates whether respondents received VDC.

Figure 2 was constructed from the variable DISABWRK, which codes the response to a question about disabilities that limit or prevent work. Men with fair or poor health were identified using the variable HEALTH, which gives self-reported health status. This variable is available beginning in 1996.

The employment measure used in Panel A of Figure 3 is based on the CPS variable EMPSTAT, which codes as working men at work, with a job not at work, or in the armed forces. The share quitting or retiring for health reasons in Panel B is based on the CPS variable QUITSICK, which identifies respondents who said that they had ever retired or left a job for health reasons.

All plots show weighted means collapsed by year using sampling weights (PERWT) and including imputed values. Because income amounts refer to the previous year in the March CPS, Figure 1 runs from 1989-2005. Disability and health are measured at the time of the survey, so Panel A of Figure 2 runs from 1990-2006 and Panel B from 1996-2006.

2000 Census Disability Questions

We constructed disability variables from responses to the following questions:

16. Does this person have any of the following long-lasting conditions:

- **a.** Blindness, deafness, or a severe vision or hearing impairment? (Yes, No)
- **b.** A condition that substantially limits one or more basic physical activities such as walking, climbing stairs, reaching, lifting, or carrying? (Yes, No)

17. Because of a physical, mental, or emotional condition lasting 6 months or more, does this person have any difficulty in doing any of the following activities:

- **a.** Learning, remembering, or concentrating? (Yes, No)
- **b.** Dressing, bathing, or getting around inside the home? (Yes, No)
- **c.** (Answer if this person is 16 YEARS OLD OR OVER.) Going outside the home alone to shop or visit a doctor' s office? (Yes, No)
- **d.** (Answer if this person is 16 YEARS OLD OR OVER.) Working at a job or business? (Yes, No)

A respondent was coded as having a work-limiting disability if he or she answered "Yes" to 17(d). He was coded as having a non-work-limiting disability if he answered "No" to 17(d), but yes to any of 16(a), 16(b), or 17(a)-17(c). The variable "any disability" was defined as having either a work-limiting or non-work-limiting disability. Specific disabilities were coded as follows: Vision or hearing (yes to 16(a)); Physical (yes to 16(b)); Mental (yes to 17(a)); Self-care (yes to 17(b)); Mobility (yes to 17(c)).

2000 Census VA and Social Security Income Questions

The 2000 Census has one multi-part question that collects information on income by source. We use these parts of question 31, Income in 1999:

d. Social Security or Railroad Retirement

e. Supplemental Security Income (SSI)

h. Any other sources of income received regularly such as Veterans' (VA) payments, unemployment compensation, child support, or alimony - Do NOT include lump-sum payments such as money from an inheritance or sale of a home.

The response to 31(h) is used to code an indicator for Other Income (mostly VDC); the response to 31(d) is used to code an indicator for Social Security Income (mostly SSDI); our dummy for any federal transfers indicates individuals with an amount in either 31(d), 31(e), or 31(h).

1987 Survey of Veterans (SOV-III)

Sample selection

The analysis of the SOV-III in Section 3.5 starts with the extract of 3,337 Vietnam and later-era veterans used by Angrist (1993). These data are available at http://econ-www.mit.edu/faculty/angrist/dat: For confidentiality reasons, age data in the SOV-III are bracketed in 5-year intervals. The sample used in Table 7, Panel A is arrived at by restricting to white males with bracketed ages between 30 and 44 who served during the Vietnam era as indicated by their response to Question 4(c). Finally, the five observations with missing education (coded N/A) are deleted to arrive at a sample size of 1893. The sample in Panel B is further restricted to recoded ages between 35 and 39 for a sample size of 724.

Variable definitions

The measures of combat exposure were taken from questions 16(b), 17, and 18:

16b. In which of these places did you serve, sail in, or fly missions over while on active duty in the United States Armed Forces? (*list of regions*)

17. During your military service, were you ever in or exposed to combat? (Yes, No)

18. Even though you were not in combat, were you ever stationed in a war zone? (Yes, No)

The dependent variables in columns (1) to (3) of Table 7 were respectively defined as an indicator for Vietnam, Laos, or Cambodia on question 16(b); and indicator for "Yes" on 17; an indicator for "Yes" on 17 or "Yes" on 18.

The measure of service-connected disability status is taken from question 35(b):

35b. Have you ever been notified by the VA that you have a medical condition or disability related to your military service or that you are eligible for VA medical care because you have a medical condition or disability related to your military experience? (Yes, No)

Schooling Group Definitions in the 2000 Census and SOV-III

The analysis in Section 3.4.2 defines four schooling groups using the highest level of school completed (question 9 in the census). The four census schooling groups are:

1. *High school dropout*: the highest level completed is at most 12th grade, no diploma (response code less than or equal to 9)
2. *High school graduate*: the highest level completed is high school graduate or GED (response code equal to 10)
3. *Some college*: the highest level completed is greater than high school graduate but less than a Bachelor' s degree (response codes 11-13)
4. *College graduate*: the highest level completed is greater than or equal to a Bachelor' s degree (response codes 14 and up)

The four schooling groups in the Survey of Veterans (analyzed in Section 3.5) are based on respondents' answers to questions 15(b) (highest degree before service), 15(c) (highest grade completed), and 15(d) (highest degree received):

1. *High school dropout*: the highest level completed is at most 11 (response code 5 or less in question 15(c))
2. *High school graduate*: the highest level completed is 12 or a vocational program (response code 6 or 7 in question 15(c))
3. *Some college*: the highest level completed includes at least one year of college but less than a Bachelor' s degree (response code 8 or more in question 15(c), but a code of less than 2 in questions 15(b) and 15(d))
4. *College graduate*: the highest level completed is at least a Bachelor' s degree (response code 2 or more in question 15(b) or 15(d))

Appendix Tables

Table A1. Proportion Draft-Eligible in the 2000 Census

	Actual			Theoretical	Difference
	HS dropout (1)	HS grad or more (2)	All (3)	(4)	(5)
			(A) White		
1948	0.5241	0.5310	0.5303	0.5348	-.0044
	(.0039)	(.0013)	(.0012)		(.0012)
1949	0.5394	0.5356	0.5359	0.5358	.0002
	(.0039)	(.0013)	(.0012)		(.0012)
1950	0.5279	0.5394	0.5384	0.5376	.0008
	(.0040)	(.0013)	(.0012)		(.0012)
1951	0.3336	0.3401	0.3395	0.3424	-.0030
	(.0037)	(.0012)	(.0011)		(.0011)
1952	0.2499	0.2605	0.2596	0.2596	.0000
	(.0034)	(.0011)	(.0010)		(.0010)
1953	0.2575	0.2594	0.2592	0.2604	-.0012
	(.0033)	(.0011)	(.0010)		(.0010)
$F(6,\infty)$	4.77	2.85	3.74	-	3.74
p-value	0.0001	0.0090	0.0010		0.0010
N	109,754	1,031,797	1,141,551		
			(B) Nonwhite		
1948	0.5413	0.5304	0.5376	0.5358	.0017
	(.0068)	(.0039)	(.0034)		(.0034)
1949	0.5294	0.5395	0.5371	0.5354	.0016
	(.0067)	(.0037)	(.0032)		(.0032)
1950	0.5365	0.5463	0.5440	0.5371	.0069**
	(.0066)	(.0036)	(.0032)		(.0032)
1951	0.3362	0.3446	0.3427	0.3417	.0010
	(.0065)	(.0035)	(.0031)		(.0031)
1952	0.2598	0.2666	0.2650	0.2583	.0068**
	(.0058)	(.0032)	(.0028)		(.0028)
1953	0.2604	0.2667	0.2653	0.2600	.0053*
	(.0058)	(.0031)	(.0028)		(.0028)
$F(6,\infty)$	0.38	3.31	2.48	-	2.48
p-value	0.8950	0.0029	0.0210		0.0210
N	37,313	117,497	154,810		

Notes: Columns 1-2 report the fraction draft eligible observed in each cohort by high school graduation status. Column 3 reports the overall fraction draft-eligible. Column 4 reports the theoretical fraction assuming births are evenly distributed within a month. Column 5 reports the difference between the overall empirical proportion draft eligible and the theoretical fraction, with robust standard errors in parentheses. The F-statistic is for a joint test of theoretical and empirical equality for all cohorts.

Table A2. Descriptive statistics by race, veteran status, and wage percentile

	Wage index percentile							
	<10		10-25		25-75		>75	
	Vietnam Vets (1)	Non-Vets (2)	Vietnam Vets (3)	Non-Vets (4)	Vietnam Vets (5)	Non-Vets (6)	Vietnam Vets (7)	Non-Vets (8)
A. Whites								
Draft eligibility (by RSN)	0.520	0.393	0.537	0.361	0.550	0.381	0.588	0.404
Age	49.4	49.0	49.5	48.8	49.6	49.0	49.8	49.2
i. Disability variables								
Any disability	0.346	0.400	0.259	0.268	0.208	0.169	0.146	0.103
Work-limiting disability	0.125	0.252	0.168	0.174	0.126	0.106	0.087	0.069
Non-work-limiting disability	0.131	0.147	0.091	0.094	0.082	0.063	0.059	0.038
ii. Transfer income								
Other income (mostly VDC) > 0	0.118	0.152	0.107	0.046	0.111	0.036	0.103	0.027
SSA income excluding SSI (mostly SSDI) > 0	0.077	0.104	0.047	0.053	0.032	0.026	0.019	0.013
SSI > 0	0.033	0.069	0.017	0.029	0.011	0.012	0.006	0.005
Any Federal transfer income > 0	0.197	0.198	0.151	0.117	0.138	0.069	0.119	0.042
iii. Specific disability types								
Mental (difficulty learning, remembering, or concentrating)	0.099	0.146	0.060	0.064	0.044	0.031	0.028	0.016
Vision or hearing (blindness, deafness, or a severe vision or hearing impairment)	0.069	0.079	0.047	0.050	0.042	0.032	0.030	0.019
Physical (limitation to physical activities such as walking, climbing stairs, reaching, lifting, or carrying)	0.178	0.198	0.122	0.116	0.096	0.069	0.062	0.035
Mobility (difficulty going outside the home alone)	0.103	0.144	0.071	0.080	0.049	0.041	0.030	0.021
Self-care (difficulty dressing, bathing, or getting around inside the home)	0.045	0.067	0.027	0.032	0.020	0.016	0.014	0.008
iv. Labor market variables								
Not working	0.292	0.339	0.200	0.208	0.141	0.118	0.090	0.069
Not in labor force	0.249	0.293	0.164	0.173	0.114	0.094	0.070	0.052
N	29,851	84,736	60,671	110,788	193,623	378,708	64543	218,650
B. Nonwhites								
Draft eligibility (by RSN)	0.492	0.415	0.493	0.407	0.520	0.404	0.537	0.406
Age	49.1	49.0	49.3	49.0	49.5	49.0	49.6	49.1
i. Disability variables								
Any disability	0.427	0.451	0.386	0.415	0.339	0.337	0.258	0.210
Work-limiting disability	0.265	0.277	0.235	0.257	0.214	0.219	0.167	0.143
Non-work-limiting disability	0.162	0.174	0.151	0.157	0.125	0.118	0.091	0.067
ii. Transfer income								
Other income (mostly VDC) > 0	0.125	0.057	0.124	0.055	0.139	0.054	0.150	0.043
SSA income excluding SSI (mostly SSDI) > 0	0.082	0.095	0.084	0.085	0.058	0.060	0.037	0.030
SSI > 0	0.062	0.090	0.050	0.077	0.029	0.047	0.014	0.019
Any Federal transfer income > 0	0.233	0.216	0.225	0.195	0.198	0.146	0.181	0.084
iii. Specific disability types								
Mental (difficulty learning, remembering, or concentrating)	0.124	0.138	0.103	0.118	0.073	0.070	0.048	0.030
Vision or hearing (blindness, deafness, or a severe vision or hearing impairment)	0.067	0.070	0.058	0.065	0.048	0.047	0.039	0.029
Physical (limitation to physical activities such as walking, climbing stairs, reaching, lifting, or carrying)	0.211	0.203	0.173	0.179	0.150	0.135	0.111	0.073
Mobility (difficulty going outside the home alone)	0.156	0.176	0.141	0.168	0.118	0.122	0.082	0.068
Self-care (difficulty dressing, bathing, or getting around inside the home)	0.072	0.068	0.049	0.064	0.036	0.042	0.026	0.021
iv. Labor market variables								
Not working	0.480	0.513	0.417	0.480	0.306	0.360	0.195	0.198
Not in labor force	0.409	0.445	0.358	0.410	0.254	0.302	0.160	0.142
N	2,365	13,637	4,518	18,630	25, 963	51,125	12579	25,993

Notes: Same as the previous table.

Table A3. Estimated Median of the VDC Replacement Rates, by States and by Skill Groups

		Replacement rate	Average annual earnings
		(A) White	
Wage index percentile	<10	.983	34556
	10-25	.848	38412
	25-75	.707	49458
	>75	.569	65411
Schooling groups	HS dropout	1.001	33382
	HS grad	.828	39636
	Some college	.730	46604
	College degree	.586	63266
		(B) Nonwhite	
Wage index percentile	<10	1.226	28859
	10-25	1.107	32699
	25-75	.932	35777
	>75	.731	47308
Schooling groups	HS dropout	1.250	29483
	HS grad	1.001	32905
	Some college	.854	38036
	College degree	.714	48266

Note: All states are included, except for DC. N=279,999 for whites and 31,352 for nonwhites. The replacement rate is the reciprocal of the OLS estimate of the coefficient of the state average VDC, for 100% disability or IU, in a regression of annual earnings, without covariates or a constant. The state average VDC with 100% disability/IU is derived from the "Review of State Variances in VA Disability Compensation Payment" (Department of VA Office of Inspector General, 2004).

Table A4. 2SLS estimates of veteran effects by predicted wage and schooling: Non-whites

		Disability variables			Transfer income			Labor Force Status	
		Any disability	Work-limiting disability	Non-work-limiting disability	Other income (mostly VDC) >0	SSA income (mostly SSDI) >0	Any Federal transfer income >0	Not working	Not in labor force
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
					A. By wage index percentile				
Veteran status x wage index percentile	<10	-.102	-.045	-.057	.045	.026	.359*	.304	.333
		(.235)	(.211)	(.177)	(.120)	(.136)	(.205)	(.242)	(.244)
	10-25	.007	.004	.003	-.101	-.026	-.029	.114	.254*
		(.144)	(.128)	(.107)	(.075)	(.083)	(.117)	(.145)	(.147)
	25-75	-.045	-.035	-.011	.046*	-.034	.028	-.051	-.026
		(.045)	(.039)	(.031)	(.025)	(.023)	(.035)	(.045)	(.044)
	>75	-.094	-.072	-.022	.049*	-.023	.007	.011	.010
		(.046)	(.039)	(.029)	(.029)	(.019)	(.035)	(.043)	(.040)
					B. By schooling group				
Veteran status x schooling group	HS dropout	-.255	-.175	-.080	-.077	-.053	.225	.227	.332
		(.236)	(.209)	(.179)	(.121)	(.141)	(.200)	(.234)	(.240)
	HS grad	-.022	-.027	.005	.016	-.031	.009	-.043	.012
		(.057)	(.050)	(.039)	(.030)	(.029)	(.044)	(.058)	(.056)
	Some college	-.009	.010	-.019	.050*	-.017	.031	.002	.015
		(.047)	(.041)	(.032)	(.029)	(.022)	(.037)	(.047)	(.045)
	College degree	-.081	-.079	-.003	.053*	-.017	.038	.063	.055
		(.048)	(.041)	(.030)	(.030)	(.019)	(.036)	(.044)	(.041)

Notes: The same as Table 5. The sample size is 154,810.

Table A5. 2SLS estimates of veteran effect interactions for specific disability types: Non-whites

		Disability Type:				
		Mental (1)	Vision or hearing (2)	Physical (3)	Mobility (4)	Self-care (5)
		A. By wage index percentile				
Veteran status x wage index percentile	<10	-.092	.177	-.144	.035	.060
		(.162)	(.123)	(.192)	(.179)	(.120)
	10-25	.085	.158**	-.053	.016	.052
		(.098)	(.075)	(.113)	(.110)	(.074)
	25-75	.023	.032	-.012	-.011	-.002
		(.025)	(.020)	(.033)	(.031)	(.019)
	>75	-.004	-.002	-.028	-.013	.015
		(.021)	(.019)	(.031)	(.029)	(.017)
		B. By schooling group				
Veteran status x schooling group	HS dropout	-.013	.258**	-.158	-.042	.095
		(.166)	(.126)	(.190)	(.182)	(.125)
	HS grad	.011	.053**	.012	-.042	-.007
		(.032)	(.026)	(.042)	(.040)	(.025)
	Some college	.049*	.024	-.037	.046	.012
		(.025)	(.021)	(.035)	(.032)	(.019)
	College degree	.006	.007	-.012	-.006	.023
		(.021)	(.020)	(.032)	(.029)	(.017)

Notes: The same as table 6. The sample size is 154,810.