Title

Essays in Labor Economics and Econometric Methods

Author

Johnston, Carla

Publication Date

2020

Peer reviewed|Thesis/dissertation

Essays in Labor Economics and Econometric Methods

by

Carla Johnston

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David E. Card, Chair
Professor Patrick M. Kline
Associate Professor Reed Walker

Spring 2020

Essays in Labor Economics and Econometric Methods

Abstract

Essays in Labor Economics and Econometric Methods

by

Carla Johnston

Doctor of Philosophy in Economics

University of California, Berkeley

Professor David E. Card, Chair

The impact of a subsidy or transfer depends greatly on the market forces surrounding the economic actors who receive them. This dissertation contributes to the field of labor economics by shedding light on two under-studied, but widely dispersed, government subsidies. These subsidies are the tax exemptions given to nonprofit firms and the Earned Income Tax Credit (EITC) issued to low-income households. Although one of these subsidies is given to a firm and the other to a consumer, both aim to alleviate the low capital burdens for underserved populations. The first chapter examines the effect of a firm's nonprofit tax status on workers' wages using the universe of quarterly wage records from the state of Florida. The second chapter investigates market effects for a large durable good during the issuance months of the EITC. The third chapter presents a theoretical model that improves partial effects estimation for ordered outcome variables, which are often variables of interest in empirical questions of this sphere. This is particularly the case when using publicly available, binned data.

In the first chapter I explore the nonprofit earnings penalty. To separate the influence of demand and supply, I leverage workers who change employers in administrative tax data. The average nonprofit worker earns 5.5 percent less than the average for-profit worker. Supply-side factors (worker selection) contribute 80 percent of the nonprofit differential. The remaining 20 percent is from demand (a nonprofit penalty). Within-worker nonprofit variation generates several insights about the influence of nonprofits on the labor market. Nonprofits compress the wage distribution and reduce inequality among earners. Nonprofit penalties are much more pronounced in classic charities than in "commercial" nonprofits, which sometimes exhibit nonprofit premia. This study is the first to harness administrative wage records, rather than survey data, to estimate the nonprofit earnings penalty. This study is also the first to show serious misreporting issues in survey data which could bias results.

The second chapter of my dissertation uses Texas DMV vehicle registration records to precisely estimate quantity and price effects for used car sales in the weeks following the issuance of the EITC. I use a difference in difference framework taking advantage of the

fact that most EITC returns are issued in February and March, and that the share of the population receiving the EITC varies greatly across zipcodes. I find that for a zipcode with the average EITC population share (20 percent), used car sales increase by 33.3 percent. I find little movement in prices. I conclude that the timing of the EITC return does not result in a large loss of incidence for the consumer, despite the large increase in demand for used cars. This is the first study to examine price effects, in addition to consumption effects, of the EITC.

My first two chapters address wage and price effects of subsidies using empirical applications. My third chapter presents a theoretical advancement in methods used to recover partial effects that are of interest in this field of economics. Often outcome variables of interest such as income, are binned, resulting in "ordered" outcomes. I present an ordered response model that relaxes rigid distributional assumptions, allowing for less biased and more consistent estimators. This model allows for flexibility gains while still being tractable and easy to implement for the applied researcher.

To Daniel and Charlotte

Contents

List of Figures

List of Tables

Acknowledgments

This dissertation is the symbolic end to a long journey and I'm grateful to all those who have enriched my life along the way. Many thanks to my adviser and my committee for helping me through the graduate school process. Reed Walker taught me invaluable research skills during my time as his RA. Enrico Moretti and David Card oversaw very constructive labor "therapy" sessions during my third year at Berkeley. Thank you as well to Josh Goldstein in the Demography Department, who supported me as I came back from maternity leave to GSI his class. While writing my dissertation I had the support of the National Institute of Health, NIA grant T32-AG000246 and UC Berkeley scholarships.

Spending four years in a windowless office together was made sunny by my cohort-mates Francis Wong, Jonathan Holmes, Evan Rose, and Maxim Massenkoff. Eating lunch with them was my favorite part of the day. Ingrid Haegele was my constant cheerleader, buoying me up when I felt discouraged and reminding me of the many talents I had to offer. My lunch talks with Katarina Jensen this past year sustained me as I struggled with the heavy load life gave me. Her compassion and friendship is something I will cherish for life. I will always look forward to frozen yogurt runs and Celia lunches with Maria Kurakina. She is a true friend and I look forward to making many more wonderful memories together.

Just as raising a child takes a village, raising a child and writing a dissertation takes many villages. I am so indebted to my many wonderful communities who have nourished me and my family. I love the bonds I have forged with the other econ parents, lamenting over sleepless nights, poopy diapers, and referee reports. Thank you Daniel Haanwinckel, Thais Vilela, and Isabela Haanwinckel, Matthias Hoelzein, Heather Bartlett, and Vera Hoelzein, Dario Tortarolo, Malena Arcidíancono, and Benjamin Tortarolo, and Peter Jones, Gabriela Wyatt, and Nolan Jones for sharing your families' journeys with me. My dream team Lori Miller, Laura Brignone, and Alicia Harper deserve outstanding recognition. Afternoon playdates and after-bedtime movie marathons made many frustrating days happier. Thank you Alison Maeser for being a true listener and Sarah Harrison for having my back. My congregation in the Church of Jesus Christ of Latter-Day Saints also provided so much support along the way, helping me grow not just academically, but spiritually. The lessons I have learned there have given deeper meaning to all my life efforts.

I started down this economics journey nine years ago when Lars Lefgren offered me a position as a teaching assistant for his microeconomics class. His support, along with the support of the BYU Economics Department, has been vital every step of the way. Thank you to James McDonald for being an excellent mentor and "life-coach." Remember, econometrics is easy, friendly and fun! Joe Price and Olga Stoddard have been immensely encouraging and a joy to reunite with whenever the opportunity affords. Working with co-authors has made my life rich with friendships and research. My first chapter is co-authored with BYU alum and long-lost cousin Andrew Johnston. My second chapter is a collaboration between Riley Wilson, my former econometrics TA, and Josh Witter. My third chapter is joint-work with James McDonald and Kramer Quist. Thank you to all of them for sharing their insights and talents. Jeff Sorensen provided invaluable guidance as I made the transition from academia to private industry.

The love and support of a family is my most cherished gift of this journey. My parents, Jeff and Nelly Johnston went above and beyond by the parental call-of-duty, moving from Salt Lake City to Berkeley to provide childcare for Charlotte during the last half of my PhD. My mother-in-law, Cynthia Gardner, has also been a huge boon to my little family. Rebecca and Kevin Gardner, and Ben and Alysha Gardner have always been a joy to have in my life. Thank you to my amazing brothers, Paul and Jeff Johnston, and to my sister-in-law, Dru Johnston. Thank you to my darling daughter, Charlotte. I am so glad to have had my life utterly transformed by you. And lastly, thank you to my husband, Daniel Gardner. We wed the week we graduated from college and have been together on this journey every step of the way. The ways he has helped me are innumerable, from dropping off dinner during stressful exam weeks, to delaying his program for a semester when we had too much to juggle. I am grateful for all the ways we have grown individually and as a couple and I'm so lucky to have him as a life partner.

Introduction

The effect of tax subsidies on market outcomes has intrigued economists for many centuries. In recent US history, government subsidies of various forms have steadily increased. The goal of these various subsidies are widespread; some subsidies aim to stimulate business activity, other subsidies aim to relieve individual liquidity constraints. Many have environmental considerations as their inspiration.

Subsidies are useful insofar as they correct ideal market aberrations. For example, if society could benefit from a certain type of business activity, but the costs of running such a business are lower than the potential profits, a government tax break could act as an incentive for individuals to create such businesses. This is the idea behind a tax break for nonprofit firms. Subsidies directed at individuals often address liquidity constraints. An individual may be prevented from making necessary purchases because they cannot access liquidity markets very easily; it might be too costly for a struggling family to search out a reasonable loan rate, for example. Additionally, markets for some types of loans come with exorbitantly high interest rates. Thus a government subsidy in the form of a tax credit can give an individual access to ready to use cash, sidestepping the problems of loan market scarcity and loan market search.

Once a subsidy is issued, policymakers have a myriad of questions regarding the effects of the subsidy. Did the subsidy make it to its intended target or did others outside of the issuance criteria manage to get a piece of it? Did the subsidy have its intended effect on its intended target? Did this subsidy have any effect on economic actors connected with the target? Are there any unexpected adverse effects? Are there any unexpected benefits? Subsidies do not exist in a vacuum. Where market structures permanently altered from the subsidy, or only altered in the short term?

Thankfully many government agencies provide the data needed to answer these questions. The IRS publishes aggregated data on many tax outcomes, such as the number of individuals who receive the Earned Income Tax Credit with income between 30,000 and 60,000 dollars. Nonprofit firms are required by law to publish their tax records. Welfare agencies also have a wide collection of publicly available data. Agencies are often able to release data to the public because they aggregate or anonymize statistics that would otherwise uniquely identify individuals. One method of aggregation is binning data, a very common practice with income. Rather than reporting exact dollar income, agencies report an individual's income as between two levels. Consistently estimated causal effects when the outcome variable is binned requires statistical tools above and beyond the standard linear regression.

This dissertation addresses these two facets of subsidy research; obtaining actual economic estimates and refining the tools used to obtain those estimates. My first chapter focuses on the tax break given to nonprofit firms. I focus on the questions of which businesses receive the tax break and how this tax break affects the workers' wages of nonprofit firms. My second chapter covers the Earned Income Tax Credit's effect on the purchasing behavior of its recipients. It also analyzes whether the large durable goods market is altered during the month of the EITC's issuance. In my third chapter I present an ordered response model that relaxes the distributional assumptions of typical models used with binned outcome data. My model allows for flexibility gains while still being straightforward to implement for the applied researcher.

Chapter 1

Is Compassion a Good Career Move?: Nonprofit Earning Differentials from Job Changes

1.1 Introduction

Over the past half century, nonprofit organizations have proliferated in number, revenue, and employment (Leete, 2001). In the past twenty years alone, the share of all workers employed by the nonprofit sector has increased by 40 percent (Friesenhahn 2016; Hirsh, MacPherson, and Preston 2017). The shift toward nonprofit employers may have consequences for the labor market if nonprofit firms affect the distribution of worker earnings (Rose-Ackerman, 1996). Nonprofits may, on one hand, pay more because they must re-invest net earnings within the organization, encouraging the firm to distribute earnings internally in the form of higher wages (Pauly and Redisch, 1973; Bishow and Monaco, 2016); on the other hand, nonprofits can reduce wages if workers derive utility from participating in the mission of the nonprofit, eliciting a labor donation (Hansmann, 1980; Preston, 1989; Frank, 1996). We evaluate these hypotheses by decomposing the nonprofit pay gap into demand- and supply-side factors.

Disentangling supply and demand in this setting is empirically challenging. Workers, for one thing, are not randomly assigned to employers. Even if random assignment were possible, the wage data used in previous studies are self-reported and contain considerable measurement error in earnings (Bound and Krueger, 1991). Using administrative data, we demonstrate that these same records also have significant measurement error in nonprofit status, introducing bias that is hard to characterize, let alone quantify.

In this chapter, we address these challenges by bringing to bear full-population earnings and tax records from Florida. By focusing on workers who transition between for-profit and nonprofit employers, we account for unobserved, worker-specific traits to decouple the role of the supply- and demand-side factors driving nonprofit earning differences. The administrative data we use cover the full working population of Florida and, because the data are derived from tax records, there are strong incentives for wages and nonprofit status to be recorded accurately.

The data reveal that nonprofits pay 5.5 percent less, on average. 80 percent of this differential is explained by worker selection, and the remaining 20 is explained by a nonprofit penalty. While the average nonprofit penalty is slight, just one percent, the penalty is much larger for high earners. The nonprofit penalty at the 95th percentile of the earnings distribution is 10 percent, ten times larger than average. This significant penalty may be the result of competitive labor-market forces in which nonprofit managers accept lower pay for greater influence over the direction of nonprofits (Glaeser, 2002). Another possibility is the influence of regulations which sanction highly paid nonprofit managers and the boards that offer compensation eventually deemed "unreasonable."

Not all workers suffer a nonprofit penalty. Nonprofits pay a premium to workers in the bottom 25 percent of the earnings distribution, suggesting that nonprofits compress wages. If one applied the earnings compression we observe in nonprofits to the for-profit distribution, it would reduce income inequality, as measured by Gini coefficients, by 60 percent.

Several papers have estimated nonprofit penalties for individual industries (Borjas, Frech III and Ginsburg, 1983; Weisbrod, 1975; Goddeeris, 1988; Preston, 1989; Holtmann and Idson, 1993; Leete, 2001; Mocan and Tekin, 2003; Hirsch, Macpherson and Preston, 2017). We shed light on industry-specific nonprofit penalties, first, by presenting visual evidence that features the earning dynamics of workers transitioning between for-profit and nonprofit work in each industry. The estimates from this event-study approach demonstrate that the nonprofit penalty varies significantly from industry to industry. Workers face the most significant penalties when working in classic charitable organizations like legal aid (–13%) and religious employers (–10%). A few industries exhibit no nonprofit differential, including hospitals and nursing homes. In some industries, workers earn more in a nonprofit than in a for-profit, including in family services (3%), outpatient healthcare (4%), and childcare centers (5%), consistent with evidence suggesting nonprofit premia in some settings (Leete, 2001; Bishow and Monaco, 2016). We explore several industry-level explanations for varying penalties. Nonprofit penalties/premia are most strongly related to differences in worker fixed effects across nonprofit/for-profit sectors within industry, suggesting again the egalitarian influence of nonprofits on the distribution of wages. We find no evidence that the nonprofit wage differences across industries are related to differences in the competitive environment, employee misattribution of nonprofit status, or industry-specific differences in nonprofit utility.

It's useful to return to the broad misclassification of nonprofits in survey records to notice what it implies. That many employees do not know the nonprofit status of their employer seems to undermine a primary explanation for nonprofit existence: nonprofit legal status allows entrepreneurs to commit to providing quality and, thereby, gain market share. But if employees don't know that a firm is nonprofit, it's hard to imagine customers do. This suggests that nonprofit status is an information signal usually intended for deliberately informed donors, rather than paying customers or employees.

This chapter contributes to a long literature investigating the economic behavior of nonprofits (Arrow, 1978; Newhouse, 1970; Feldstein, 1971; Baumol and Bowen, 1965; Horwitz and Nichols, 2007). We show that the survey data used to study this question in previous research contain significant measurement error in nonprofit designation (e.g., at least half as many workers misclassify their status as there are nonprofit workers, greater than 4 percent of all respondents). This chapter is the first to resolve this issue using full-population, administrative panel data to account for individual worker differences and illuminate the magnitude of the nonprofit wage penalty in various settings. The size and scope of the data allow us to leverage a design-based approach to answer the question while providing clear visual evidence in event-study figures.

Our work compares most closely to Ruhm and Borkoski (2003), and later Hirsch et al. (2017), who use the Outgoing Rotation Group of the Current Population Survey to study workers who transition to or from nonprofit settings in survey data providing two observations, one year apart. Our primary contribution relative to these studies is that we (1) leverage administrative tax data, significantly reducing the scope for mismeasurement in both earnings and nonprofit status; (2) study long panels of individuals changing jobs to carefully account for job-change dynamics; and (3) harness the experience of several tens of thousands of workers who transitioned between nonprofit and for-profit employment to provide statistical clarity.

1.2 Background

To avoid a contradiction in terms, what is called "profit" in a typical setting is called "net earnings" in a nonprofit organization (revenues less cost). The essential characteristic of a nonprofit is that the organization is barred from distributing earnings to owners or managers, an institutional rule described by economists as the "non-distribution constraint" (Hansmann, 1980).

The primary economic rationale for the institutional feature is to mitigate concerns arising from information asymmetry. Should Jane donate money to charity, she cannot easily verify whether promised services were furnished to the indigent. If the charity were organized as a for-profit firm, its owner would be tempted to withhold promised services for personal gain. The non-distribution constraint blunts this incentive, allowing Jane to have greater confidence that her donation reaches the intended beneficiary. Similar information asymmetries exist in personal services (like assisted-living facilities, hospitals, day cares, and schools) in which the quality of care cannot easily be assessed by the patron. In many cases, the service recipient is unhelpful even in evaluating quality since the beneficiaries may be sedated, disabled, a child, or otherwise unable to determine the quality of care due to its technical nature, as is often the case when consumers seek medical treatment.

Jane can have confidence that the penalties imposed for violating the non-distributional constraint are quite exacting. Board members that approve a compensation package eventually deemed "unreasonable" [1] by the IRS are required to pay a fine equal to ten percent

[1] The classification of a compensation package as "unreasonable" is somewhat subjective and determined

of the overage (IRS 2016). In addition, the (overpaid) manager must repay the overage to the nonprofit, including interest, in addition to paying a 25 percent excise tax on the overpayment (Ibid). Under the uncertainty of this somewhat subjective rule, board members and managers may agree to lower levels of compensation to avoid censure and fine, and potentially find nonpecuniary avenues to transfer utility. As an aside, this is one possible explanation for the sizeable nonprofit wage penalty we discover among the top percentiles of the wage distribution.

In exchange for the non-distribution constraint, the US government grants nonprofit organizations an exemption from federal income taxation under the US Internal Revenue Code section 501(c). Entrepreneurs can incorporate their organizations as nonprofits if they fit into one of several categories: traditional charities, religious communities, scientific organizations, education providers, and organizations that work to prevent child cruelty (section 501(c)3). Donations to these groups are tax deductible [2]. Nonprofit employees pay individual income taxes on their earnings, as they would if they were employed in for-profit institutions. Nonprofit employers are liable for payroll taxes that fund social insurance programs, but they do not pay federal or state income tax and do not pay property taxes—true in all 50 states (Lindblad, 2199). In our setting in Florida, nonprofits are also exempt from paying sales and use taxes, but this is not everywhere true.

1.3 Data

Measurement Error of Nonprofit Status in Survey Data

Accurately gauging nonprofit differentials depends on reliable measures of nonprofit status. It is well known that survey data contain considerable measurement error in self-reported earnings arising from rounding, seam bias, imperfect memories, and intentional misrepresentation (Bound and Krueger, 1991), in addition to selective reporting and top-coded earnings (Hirsch et al., 2017) [3]. What has been unexplored is whether respondents accurately identify the nonprofit status of their employer when completing surveys like the Current Population Survey (CPS) or the American Community Survey (ACS). On one hand, an employer's nonprofit status is binary and stable, so it seems reasonable that employees may be able to reliably recall nonprofit status. On the other, employers may have little reason to communicate their tax status with workers.

To assess the prevalence of measurement error, we compare the nonprofit attribution in the ACS coverage of Florida with administrative employment records covering the same state. We reveal high rates of misidentification. In table 1.1, we compare the nonprofit

by the IRS.

[2] Hospitals often enjoy charitable/nonprofit status. This is a holdover from an era in which hospitals were charities that provided health services to the indigent (Hansmann, 1980)

[3] For instance, about 30 percent of working respondents in the CPS do not report their earnings (Hirsh, MacPhereson, and Preston 2017).

employment share in survey data to the nonprofit employment share in administrative records for several industries, focusing on those that have large nonprofit representation. In the administrative data, 72 percent of healthcare workers are employed at a nonprofit hospital; in survey records, however, only 43 percent of workers report working for a nonprofit, implying a misidentification rate of at least 40 percent. In the education sector, employees tend to make the opposite error: more than 12 percent of for-profit employees incorrectly respond that they work for nonprofits. These misidentification rates could be far higher since we are only able to ascertain net mismeasurement, not gross. For instance, should two individuals make opposite errors identifying their employers' tax status, we will detect no (net) measurement error, despite the fact that the nonprofit status of neither is correct.

The measurement problem poses difficulty for consistent estimation from survey records. From our administrative records, we can calculate a lower-bound for measurement error by summing the net error in each industry. We find that at least half as many workers as there are nonprofit employees misidentify their nonprofit status in the ACS over this period.

Measurement error of this magnitude, in the primary independent variable of interest, has likely led to significant statistical bias in estimates (Card, 1996), a challenge addressed by the administrative tax records we use. We assess the potential bias and find that the estimates resulting from mismeasurement could either attenuate or exaggerate nonprofit differentials depending on the correlations between misreporting and income.

A careful reader may notice that this broad misidentification of nonprofit employment also poses a challenge to one compelling economic rationale for nonprofit existence. Entrepreneurs elect to originate nonprofits rather than for-profits to commit to—and signal—quality in markets where quality is important but difficult for consumers to observe (Arrow, 1978; Hansmann, 1980). At first appearance, our finding that many employees do not know the nonprofit status of their employer seems to undermine this explanation. After all, it is unlikely that customers would be better informed regarding a firm's nonprofit status than employees, since any information available to customers would, by the same avenues, also be available to workers. This suggests that nonprofit status is an information signal often intended for deliberately informed donors, rather than paying customers or employees.

Data Construction

We obtained employer-employee matched administrative data for the full population of workers and employers in Florida from 2003 to 2012, and we link two large registers using identification numbers for workers and firms. The data include total earnings at each job in every quarter for the universe of legitimate workers [4]. Because the administrative earnings records are based on firms' reports used to calculate UI tax liabilities and benefits, they are subject to audit and are thus unlikely to contain significant measurement problems. Moreover, whereas survey data give rise to measurement error in the primary independent variable

[4] The data covers all businesses, nonprofit organizations, state or local government employers, and Indian tribal units that either have a yearly payroll exceeding $1,500 or have at least one employee working at least a portion of one day during any 20 weeks of the year.

of interest, the records we use to code “nonprofit” capture the firm’s official legal status[5] [6].The firm identification number in the wage records allow us to link worker wages to firm information including administrative records of their nonprofit status and detailed industry codes (NAICS).

Our main analysis centers on the earning dynamics of those who transition from for-profit to nonprofit work, while accounting for the wage evolution common to workers moving between for-profit employment. To focus the analysis on relevant individuals, we limit the data to those earnings observations in which employees were working for for-profits (those with the legal classification of c-corporation or s-corporation in the employer tax data) and those working for nonprofits (those classified as not-for-profits in the tax data). To generate a panel of worker wages for each individual, we include only the highest wage record for each worker in a given quarter when a worker has multiple jobs at one time. We drop wage records in which the employee earns less than what they would earn if they were employed full-time at the minimum wage to concentrate the analysis on similar employment arrangements, similar to Song, Price, Guvenen, Bloom and Von Wachter (2019). Some workers appear to change jobs frequently. We remove work spells with fewer than six quarters, limiting the analysis to those that have at least a year and a half of work experience both before and after a job-change “ event”. Several workers present more than one event. To leverage all the available variation, we stack events so that a given worker’s wage evolution at a given employer may function as the pre job-change earnings in one event and the post job-change earnings in a separate event.

Once the records are narrowed to workers that change jobs with at least a year and a half of tenure before and after a move, the analytic sample includes 92,429 transitions to nonprofits from for-profit firms and 66,928 transitions the other way, with 18,838 individuals transitioning in both directions at different times. In total, we leverage the wage dynamics of 178,195 nonprofit/for-profit job transitions. In the primary specification, we use 1,596,220 within-sector transitions to control for the wage dynamics general to job changes. In table 1.2, we present summary statistics for average quarterly earnings in each industry by nonprofit status.

Although these data are complete and detailed, they have important limitations that bear mention. First, the analyst has no direct information with which to compare the type or difficulty of work required in each employment setting (such as hours, work requirements, or non-wage benefits), potentially missing important non-monetary compensation differences. Second, the data do not allow the researcher to see whether job changes coincide with shocks to human capital, for instance the onset of a debilitating medical condition or the occurrence of a life-changing accident. We expect these events to be uncommon and second order.

[5]Similarly, governments have a strong incentive to make sure that firms do not erroneously report their tax-exempt status.

[6]See Salamon and Sokolowski (2005) for more information on how states collect wage records.

1.4 Empirical Methodology and Results

The ideal design to measure the nonprofit earnings penalty would be to randomize workers to sectors, for-profit or nonprofit. Absent such an experiment, researchers have sought to compare workers with similar observable characteristics across sectors Preston (1989); Leete (2001). These cross-sectional designs can provide insight but are unable to fully resolve the underlying concern that nonprofit workers may be different in unobserved dimensions, principally those related to productivity. To address this fundamental issue, we adopt two primary strategies. The simplest is a within-worker comparison in which we compare a given worker's earnings at a nonprofit to their earnings at a for-profit firm using worker-level fixed effects. The second follows a generalized difference-in-difference approach which explicitly adjusts for the earnings dynamics of job changes.

Visual Evidence from Event Studies

In addition to contributing within-worker variation and administrative data, we shed new light by presenting visual evidence of the nonprofit penalty using event studies of workers who changed jobs. In each job-change event in the data, we denote $t = 0$ the quarter in which the individual begins her new job and index all other quarters relative to it. In the baseline specification, we include six quarters leading up to the job change and 12 quarters after the event. We denote W^q_{iste} the log earnings of individual i, in year-quarter s, at event time t, in the dynamics of event type q, which describes the type of employment change. There are four event types possible: for-profit-to-for-profit transitions ($P \to P$); for-profit-to-nonprofit transitions ($P \to NP$); nonprofit-to-nonprofit transitions ($NP \to NP$); or nonprofit-to-for-profit transitions ($NP \to P$). The primary event studies we present compare for-profit-to-nonprofit transitions ($P \to NP$) with for-profit-to-for-profit ones ($P \to P$) because they provide many treatment and control events (relatively few events originate from nonprofits). We run the following regression separately for each event type:

$$W^q_{iste} = \sum_{j \neq 1} \alpha_j X1[j = t] + \sum_{y} \beta_y X1[y = s] + \gamma_{i(e)} + \epsilon^q_{iste} \tag{1.1}$$

where we include a full set of event-time dummies (α), year-quarter dummies (β), and individual fixed effects (γ) which account for the average earnings of an individual. In some specifications, we include individual-event specific fixed effects which account for the worker's earnings around the time of a given event and thus controls more flexibly for evolving human capital. We omit the event-time dummy at $t = -1$, so the event-time coefficients measure earnings relative to the quarter just before a job change. By including year-quarter dummies, we control non-parametrically for time trends including those arising from the business cycle. We can identify each dummy set because there is variation in event time driven by the variation in time when a given worker changes jobs. Throughout the analysis, standard errors are clustered at the worker level which we view as a suitably conservative.

We plot the resulting αs from these models to illustrate the dynamics of job changes and present visually how nonprofit compensation differs, conditional on worker unobservables via fixed effects. In figure 1.1, we see the earnings evolution of employees who started in for-profit firms and changed jobs. The light grey evolution reflects the earning dynamics of workers transitioning from for-profit firms to another for-profit firm. This grey line provides a baseline for how we might expect earnings to evolve for workers who change jobs, but not sectors. Workers earn slightly less in the quarter they depart and the quarter they begin a new job but maintain relatively constant wages before and after the job change. The nearly 20-percent dip in earnings in the first quarter of the new job is an artifact of the quarterly nature of the data. Unless all workers begin their employment on the first day of a quarter, quarterly earnings records will reveal lower earnings at a new job since the worker registered earnings for only a part of the quarter. The fact that earnings do not increase substantially over time is the result of the control strategy in which we account for year-quarter specific fixed effects that absorb the typical time-driven increases in earnings workers experience. The event-study figure suggests nonprofits pay a modest earnings penalty which attenuates over time. Over the three-year post period, the average nonprofit penalty is 0.9 percent. Visually, workers entering nonprofits converge to the earnings of those entering for profits. In order to compare like estimates, throughout the empirical exercises we restrict the sample to those observations used in the event studies.

Estimating the Nonprofit Penalty

One concern with comparing pre- and post-change earnings is that job changes may be related to changes in roles or status that could bias estimates if, for instance, job changes tend to occur as the result of layoffs or promotion. To address this issue, we adopt a generalized difference-in-difference approach that leverages the sharp changes in sector that take place when workers leave the for-profit sector for nonprofit employment while controlling for the dynamics that exist for job-changes within the for-profit sector, essentially adapting the event studies presented in the previous subsection to produce estimates of the nonprofit penalty. This method compares the dynamics of workers transitioning to nonprofits to the natural evolution of earnings as workers change jobs within the for-profit sector. Although job changes are not exogenous, the job-change event generates a sharp change in employer that is arguably orthogonal to unobserved determinants of wage outcomes (experience, health, ability, etc.) which likely evolve smoothly over time.

To implement the generalized DiD estimate, we denote $t = 0$ the quarter in which the individual begins his new job and index all other quarters relative to it. In the baseline specification, we concentrate on quarters close to the event, including six quarters leading up to the job change and 12 quarters after the event. Denoting W_{iste} the log-earnings of an individual in year-quarter s, at event time t, as part of event e. The primary estimates we present make within-worker comparisons among those who shift between sectors while including workers who transitioned between for-profit employers as a comparison:

$$W_{iste} = \rho NP_{iste} + \sum_{j \neq 1} \alpha_j X1[j=t] + \sum_{y} \beta_y X1[y=s] + \Gamma X + \gamma_{i(e)} + \epsilon_{iste} \tag{1.2}$$

We include a full set of event-time dummies (α), year-quarter dummies (β), and, importantly, personal dummies (γ) or finer dummies designating each individual event a person engages. The coefficient on NP, ρ, captures the average nonprofit penalty. The vector X represents various controls; in the preferred specification, we include county fixed effects since nonprofit employers tend to locate in counties with higher earnings. In the main results, we present a specification that includes industry fixed effects to evaluate whether the nonprofit penalty appears primarily within, or across, industries.

Table 1.3 presents the main results. In the cross section, nonprofit workers earn 6.9 percent less than for-profit workers employed at the same time. Including county-level controls (i.e., county fixed effects) reduces this cross-sectional difference by a quarter. When we include worker fixed effects, we find that 78 percent of the cross-sectional difference is explained by worker differences (compare columns 3 and 4). Event dummies attenuate the difference slightly more than worker effects, suggesting that cross-sectional nonprofit differences are, in part, a product of life-cycle earnings differences (compare columns 4 and 5). Finally, when we include industry fixed effects (3-digit NAICS classifiers), the nonprofit penalty attenuates little, just 10 percent, suggesting that the remaining nonprofit penalty exists primarily within industry.

The results tend to suggest smaller cross-sectional differences in compensation between nonprofits and for profits than those registered in past studies. We register a 5.5 percent cross-sectional difference whereas previous studies suggest somewhat larger gaps; Preston (1989) reports differences ranging from 0 to 32 percent and Leete (2006) reports differences between 6 and 15 percent. The most similar analysis to ours, Ruhm and Borkoski (2003), finds a 11.7 percent gap. The fact that we uncover smaller cross sectional differences could be a byproduct of measurement error in nonprofit identification in past studies if, for instance, workers in low-earning jobs were more likely to believe they worked in a nonprofit either because they worked in charitable (low-paying) settings or if low-paying employers lead their staff to believe the operation is not for profit. Leete (2006) finds no nonprofit penalty when controlling for observable characteristics (nonprofit workers earned 0.1 percent less and the confidence intervals ruled out penalties larger than 0.3 percent). When using within-worker transitions, Ruhm and Borkoski (2003) find a larger penalty of 1.0 percent where the standard errors are nearly as large, creating a wide range of plausible penalties and premia. We find a similar point estimate to Ruhm and Borkowski and—thanks to the large tax data available to us—the standard errors are tight providing quite a precise estimate. Ruhm and Borkowski's confidence intervals spanned from -3.1 to 1.1; our estimates rule out over 90 percent of that interval, providing significant statistical clarity.

The Influence of Nonprofits on the Income Distribution

In addition to seeing how nonprofit employment affects earnings on average, we explore how nonprofits shape the distribution of earnings by studying heterogeneity in nonprofit penalties in various quantiles of the income distribution. When benefits or penalties are associated with firm characteristics there is often a question of which workers are receiving such benefits (Card, Cardoso and Kline, 2016). We implement specification 1.1 for 20 pre-event income ventiles to study how the nonprofit penalty varies along the income distribution. We visualize the results in figure 1.2 by plotting the estimated wage penalties at each percentile of the pre-event income distribution τ. Workers in the lowest pre-event earnings ventiles receive a 3-percent earnings premium in nonprofits. The small, positive premium declines along the pre-event (that is, before the job change) income distribution, becoming negligible at the 30th percentile. A nonprofit penalty emerges at 40th percentile which hovers near 4-percent through the 90th percentile. At the upper reaches of the income distribution, workers pay a significantly larger penalty when working in nonprofits. At the 95th percentile, for instance, the typical earnings penalty is 10 percent, an order of magnitude more than the average penalty. At the 99th percentile, the nonprofit penalty is large at 7.5 percent. In the for-profit distribution, the top 1 percent of earners earn 10.4 percent of the income. When we apply the distribution of the nonprofit penalty to the earnings distribution of for-profit workers, we find it shrinks the Gini coefficient by 60 percent. This suggests nonprofits have an egalitarian influence on the income distribution by compressing wages, especially at the high end.

Why does the nonprofit penalty take this shape along the income distribution? Especially, why do high earners face such significant penalties? One possibility is that nonprofit managers have significant discretion over the focus and direction of their organizations which may be a valuable form of nonmonetary compensation, consistent with Glaeser (2002). Related is a second explanation in which the IRS's oversight of management compensation in nonprofits may discourage nonprofits from making generous offers to managers. If so, the market could plausibly clear when taking into account other dimensions including discretion in hiring or new initiatives. Because nonprofits cannot reward owners or managers with net earnings, the presence of lower compensation at the top of the distribution is consistent with the more-than-binding influence of the non-distribution constraint.

Industry-Specific Event Studies

Nonprofits encompass both traditional charities (e.g., churches, civic organizations) and commercial enterprises (e.g., insurance companies, health providers, broadcasting networks) and these diverse types of employers may pay differently by nonprofit status. Several prior studies estimate a nonprofit penalty for a particular industry. We contribute to these industry-specific studies by plotting the earnings evolution of workers originating in the same sector of the same industry who migrated into different sectors, e.g., comparing how the earnings of for-profit workers changed when transitioning jobs to another for-profit employer as compared to those transitioning to a nonprofit employer in the same industry.

First, we will walk through a representative figure visualizing the nonprofit earnings penalty in the legal industry, seen in figure 1.3. Those transitioning to another for-profit legal employer earn slightly more relative to the last quarter of employment in their former job, capturing the dynamics typical of changing jobs. In contrast, workers transitioning from a for-profit legal employer to a nonprofit one experience a significant drop in earnings which persists over the observation period. Before the event, earnings trends between the two groups are parallel and essentially identical, suggesting similar underlying dynamics in the two groups of workers. The relative fall of those transitioning to nonprofit work reflects the fact that among those transitioning from for-profit legal work to other legal firms, nonprofit workers bear a 17.1 percent penalty (p ¡ 0.001) compared to their for-profit alternative, similar to Weisbrod (1975) who estimated a 20-percent nonprofit penalty in law. Religious employers and other classic charities do not have a significant for-profit share. To generate the control event for each of the other classic charities, we identify the three 3-digit NAICS codes of workers that most commonly transition to that particular nonprofit type and identify workers transitioning between for-profit jobs in those industries. We find significant, visible declines in earnings for those transitioning to religious employers, civic organizations, and social advocacy groups compared to those transitioning to other for-profit employers (figure 1.3). We observe convergence over time between the nonprofit earnings profile and that in for profits. We do not observe this convergence in commercial nonprofits. It may be that classic charities have additional flexibility with workers to pay them less than market rates while vetting, training, or acculturing them. It may also be that new workers in charities pay a penalty, but established workers receive market rates for their service as the worker becomes core to the function of the organization. For workers originating from for-profit employers, the earnings paths of those moving to for profits and nonprofits are predicted to converge midway through quarter one of the fourth year after the transition.

The nonprofit penalty is not as large in several other industries. In figure 1.4, we present a parallel figure for workers transitioning from for-profit education firms to either another for-profit educator or a nonprofit educator. In this setting, pre-event earnings trend in parallel and are overlapping. After the job change, workers migrating to nonprofits appear to have no systematic wage disadvantage when compared to peers moving to for profits. In some industries, like outpatient healthcare (also in figure 1.4), we observe that workers transitioning to nonprofit employers enjoy a significant wage advantage over their for-profit counterparts. Is this apparent nonprofit premium driven by workers shifting toward subindustries that are higher paid? To shed light on this, we show the event study for job changes from for-profit doctor's offices to nonprofit doctor's offices, a subset of outpatient healthcare. Here, we find similar nonprofit premia suggesting a nonprofit advantage not driven by subindustry sorting.

Nonprofit Penalties over the Business Cycle

Though nonprofits cannot distribute net earnings, they need not spend down their surplus each year which may help them weather downturns with a cushion of savings stored during expansionary years. We leverage within-worker variation to study how the nonprofit penalty varies over the business cycle. To estimate nonprofit penalties over the business cycle, we implement the following specification.

$$W_{ste} = \sum_{j=2003}^{2012} \theta_j X N P_{ste} X 1[year = j] + \sum_{y} \beta_y X 1[y = s] + \Gamma X + \gamma_e + \epsilon_{ste} \tag{1.3}$$

The θs are the coefficients of interest on the interaction of the nonprofit indicator with a year indicator. We include year-quarter dummies (β), individual-event dummies (γ), and county controls. Including industry fixed effects yields similar results. The nonprofit gap is estimated similarly but lacks individual or event FEs, and the coefficients for both are plotted in figure 1.5. Before the recession, the nonprofit penalty was similar to the total differential suggesting little difference in worker fixed effects. In 2005, for instance, the cross-sectional wage difference was 8 percent, and the nonprofit penalty was 6.5 percent. During the recession, the nonprofit gap fell slightly, while the nonprofit penalty fell to zero by 2008, and became a 2–3 percent wage premium from 2009 through 2012, potentially the result of nonprofit cash stores.

The fact that the nonprofit differential remained negative, while the nonprofit penalty shrank and became a premium during the recession suggests that nonprofits either maintained worker earnings in the recession, or kept earnings at the same level relative to for profit firms, while the composition of nonprofit workers became less skilled. The new lower-skilled workers at nonprofit jobs earned more than they would at a for-profit job, either because for profit firms in general pay lower-skilled workers less, or because for profit firms cut worker payments across the board during the recession. This would account for the reducing nonprofit penalty and emergence of a modest nonprofit premium.

To test directly whether nonprofit firms substituting towards lower-skilled workers during the recession, we estimate AKM models to recover worker fixed effects (Abowd, Kramarz and Margolis, 1999). We model log quarterly wages w_{it} of individual i in year t as a worker component α_i, a firm premium $\phi_{J(i,t)}$, and controls contained in $x'_{it}\beta$ (including year, county, imputed experience), and an error term, ϵ_{it}.

$$w_{it} = \alpha_i + \phi_{J(i,t)} + x'_{it}\beta + \epsilon_{it} \tag{1.4}$$

Following AKM, we interpret the worker effect α_i as human capital factors (such as skills, education, ability) that are rewarded equally by employers. We interpret the establishment effect $\phi_{J(i,t)}$ as a proportional pay premium or penalty that is paid by establishment j to all its employees. Using our full sample (data from 2003-2012), we recover worker and firm fixed effects. We then plot the average worker fixed effects levels for nonprofit and for-profit firms over time. Figure 1.6 demonstrates that during the recession, nonprofits substituted toward workers with lower FE at the onset of the recession, while for-profits followed their trend line. This and the evidence above are consistent with nonprofits substituting toward workers with lower worker FE, relative to for-profit employers.

Industry-Specific Estimates of Nonprofit Penalties and Premiums

The event studies show visually that nonprofit penalties vary significantly by industry. We use equation (2) to estimate the nonprofit penalty in each industry, found in table 1.4. The preferred model estimates the within-worker nonprofit wage differential ρ while accounting for the wage dynamics of workers changing jobs within sector in each industry, seen in column 3.

In table 1.4, we compare the wages of workers as they transition between nonprofit and for-profit work in various industries. For instance, some insurance carriers are nonprofit while others are for-profit. When we look at worker transitions within this class, we find that a given worker is paid 4 percent less when working for the nonprofit insurance provider. Similarly, commercial banks can be registered as corporations or as nonprofits (a nonprofit commercial bank is sometimes known as a credit union). When we examine workers transitioning to and from nonprofit banks to and from for-profit banks, we learn that those workers earn 3 percent less in the nonprofit setting. In contrast, we find that outpatient healthcare workers earn 4 percent more in nonprofits. To make sure we are comparing like settings, we condition on those nonprofit and for-profit employers that work in "physician offices" and find that a given worker earns 3 percent more in the nonprofit setting. We estimate comparable models for traditional nonprofit charities. As in the event study, our sample of for-profit workers come from industries that have a high probability of receiving or giving workers from charity-type firms. We find significant penalties associated with working for these traditional charities on the order of 13 percent for law firms. Religious bodies pay a given worker 10 percent less, civic organizations pay 5 percent less, social advocacy groups pay 2 percent less.

1.5 Discussion

In classic charities (e.g., churches, philanthropies, and civic organizations), nonprofit workers tend to take a pay cut, evidence of a labor donation. In "commercial" nonprofits (e.g., insurance providers, commercial banks, healthcare providers), however, firms pay an attenuated penalty, and, in some industries, nonprofits pay as much or more than their for-profit peers—a striking feature. Firms can only rely on a labor donation from workers if the marginal worker is willing to accept lower wages for the warm glow of an employer (Rose-Ackerman, 1996). Even if some workers would be willing to accept lower wages for employment at nonprofits in a given industry, labor markets with lots of nonprofit employment may have to raise wages to attract the marginal worker who is unaffected by warm glow (Jones, 2015); this provides a plausible explanation for the divergence in the penalties across industry, but we find no evidence, for instance, that larger nonprofit sectors in an industry exhibit smaller penalties. Given that many workers misclassify the nonprofit status of their employer in commercial industries, labor donation, δ, may not be a significant factor in those settings. We find, however, that the nonprofit penalty in each industry is not significantly correlated with the misidentification rate in that industry ($p = 0.974$).

An unobserved component of w includes nonwage benefits and amenities. If nonprofits differ in nonwage benefits or work requirements by industry, this variation could explain differing nonprofit penalties across industries. Bishow and Monaco (2016) present data suggesting that nonwage benefits are roughly proportional with wages in nonprofits in various industries. Hirsch, MacPherson, and Preston (2017) show evidence that nonprofits workers work 4 percent fewer hours than do for-profit employees, suggesting that nonprofit wage penalties could reflect a compensating differential for a less demanding work environment. To gauge unobserved nonprofit utility by industry, we calculate the length of the average employment spell for nonprofits and for profits in each industry, which we use as a measure of how happy employees are at each type of employer. Though employment spells vary in length between nonprofits and for-profits in each industry, these differences do not predict industry-specific nonprofit penalties ($p = 0.833$), suggesting that nonwage utility does not explain the earnings penalties across industries.

We find, however, that nonprofit penalties/premia are strongly related to cross-sectional differences between the earnings of nonprofits and for-profit employers. That is, when nonprofits in a given industry pay higher-income people relative to the for-profit sector, the nonprofit within-worker premium is also larger ($p < 0.01$), which is interesting but not highly indicative of any hypothesis we have considered. To be precise about this statement, when nonprofits tend to employ lower-earning workers, other workers in that industry are paid less, conditional on their worker FE. This suggests the variance in nonprofit penalties/premia could result from productivity spillovers or the egalitarian norms of nonprofit firms.

1.6 Conclusion

The literature features a debate between analysts who argue nonprofit wage differences arise from demand-side factors (that is, that nonprofits pay differently) (Weisbrod 1983; Borjas, Frech III, and Ginsburg 1983), and others who contend these differences arise from supply-side factors (namely, nonprofits employ a different kind of worker) (Goddeeris 1988; Holtmann and Idson 1993). For example, a nonprofit wage gap could arise from disproportionate labor supply of less experienced or less educated workers who also would earn less in for-profit employment (Hirsch, Macpherson, and Preston 2017). The difficulty in this debate is that nonprofit workers may differ in a host of unobservable ways that are challenging to assess. The purpose of this chapter is to evaluate the nonprofit differential using administrative data on workers who changed jobs to account for differences in worker characteristics, which persuasively controls for unobserved personal factors that might otherwise bias measures of the nonprofit penalty. Nonprofit wage penalties from the demand side suggest labor donation in classic charities, a topic of sustained interest (Hannsman 1980; Hirsh, MacPherson, and Preston 2017).

We find that lion's share of the nonprofit gap is a product of worker composition (supply-side factors) and that a small but distinct share is attributable to a nonprofit penalty (on the demand side). These penalties are large in classic charities and smaller in commercial nonprofits where nonprofit premia appear in some industries. Though we explore several explanations for the varying nonprofit penalty/premium, we cannot find a convincing explanation. Understanding why some nonprofits pay more may illuminate policies that promote greater wages among workers. These estimates build on prior literature by accounting for worker-specific unobservables in a large dataset which allows for unbiasedness, but also statistical precision which rules out more than 90 percent of the confidence intervals provided by previous work. We also find that nonprofits compress the earnings distribution, especially at the high end, suggesting that the rapid growth of the nonprofit sector may have fostered greater income equality than would have otherwise existed.

1.7 Tables and Figures

Table 1.1: Measurement Error in Nonprofit Status

	Within Industry Percent Nonprofit			
	ACS	Admin	PP Error	% of Total NP workers
Hospitals	43%	72%	-29	40%
Educational Services	43%	49%	-6	19%
Ambulatory Health Care Services	12%	14%	-2	10%
Social Assistance	49%	55%	-6	8%
Nursing and Residential Care Facilities	25%	27%	-2	7%
Religious and Civic Organizations	100%	42%	58	4%
Recreation industries	5%	12%	-7	2%
Credit and Banking	15%	7%	7	2%
Scientific and Technical Services	2%	2%	0	2%
Utilities	8%	12%	-4	1%

Note: The first column is the percentage of reported nonprofit workers in each industry from the ACS Florida sample in 2010. The second column is the percentage of recorded nonprofit workers from the universe of Florida's UI records in 2010. The third column is the percentage-point difference between column 1 and 2. The fourth column is the industry's share of all nonprofit workers according to UI records.

Table 1.2: Industry Composition and Nonprofit Earnings Differences

	Average Quarterly Earnings ($)				
	Overall	For-profits	Nonprofits	Nonprofit Differential	Share nonprofit
	(1)	(2)	(3)	(4)	(5)
All Industries	$11,394	$11,476	$10,633	0.07	0.10
Health and Human Services					
Outpatient Health Care (621)	$13,103	$13,449	$11,739	0.13	0.20
Doctor's Offices (621111)	$16,256	$16,086	$17,618	-0.10	0.11
Hospitals (622)	$11,352	$11,206	$11,413	-0.02	0.71
Nursing Facilities (623)	$7,824	$7,740	$7,900	-0.02	0.53
Social Services (624)	$7,321	$6,969	$7,505	-0.08	0.66
Childcare (62441)	$5,873	$5,438	$6,652	-0.22	0.36
Education (611)	$10,522	$10,380	$10,646	-0.03	0.53
Finance and Management					
Law Offices (54111)	$16,173	$16,235	$9,781	0.40	0.01
Banking & Credit (522)	$13,225	$13,443	$8,751	0.35	0.05
Investments (523)	$25,460	$24,567	$40,427	-0.65	0.06
Insurance (524)	$12,936	$12,938	$11,058	0.15	0.00
Administration (561)	$9,995	$9,993	$10,853	-0.09	0.00
Utilities (221)	$18,545	$18,837	$11,164	0.41	0.04
Classic Charities					
Religious Organizations (8131)	$8,541	$7,939	$8,659	-0.09	0.84
Grantmaking Foundations (8132)	$10,194	$11,084	$10,157	0.08	0.96
Social Advocacy (8133)	$8,605	$9,687	$8,403	0.13	0.84
Civic Organizations (8134)	$9,220	$9,766	$8,998	0.08	0.71

Note: Summary statistics are calculated using the sample which includes all workers for the years 2003-2012. No sample restrictions are imposed. "Share nonprofit" indicates the share of industry workers which are employed by a nonprofit firm.

Table 1.3: Nonprofit Differential Estimates

	Log Earnings					
	Cross-Sectional Difference			Within-Worker Estimates		
	(1)	(2)	(3)	(4)	(5)	(6)
Nonprofit	-0.069***	-0.053***	-0.055***	-0.012***	-0.010***	-0.009***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Year-quarter FE	X	X	X	X	X	X
County FE		X	X	X	X	X
Event-time FE			X	X	X	X
Worker FE				X		
Event FE					X	X
Industry FE						X
R-squared	0.01	0.02	0.03	0.81	0.82	0.82
Observations	26,919,859	26,919,859	26,919,859	26,919,859	26,919,859	26,919,859

Note: Table is based on the estimation of equation 1.2 where the dependent variable is log quarterly earnings. All sample restrictions described in section III are imposed. Columns 1-3 provide estimates without controlling for person fixed-effects. Columns 4 includes a worker fixed effect, while Columns 5 and 6 include fixed-effects for events, allowing a worker with multiple events a separate FE for each event. Industries are grouped by 3-digit NAICS codes. This table leverages 1,336,205 unique workers and 1,568,483 unique job-change events. (* $p \leq 0.05$; ** $p \leq 0.01$; *** $p \leq 0.001$)

Table 1.4: Nonprofit Differential Estimates by Industry

	Nonprofit Estimates	Differential		Difference Decomposition	
	Cross Sectional Difference	Within-Worker Estimate	Gen-Diff-in-Diff Estimate	Percent Demand	Percent Supply
	(1)	(2)	(3)	(5)	(6)
All Industries	-0.053***	-0.012***	-0.010***	19%	81%
	(0.001)	(0.001)	(0.001)		
	26,919,859	26,919,859	26,919,859		
Health and Human Services					
Outpatient Health Care (621)	0.068***	0.035***	0.035***	51%	49%
	(0.010)	(0.004)	(0.004)		
	905,031	1,030,102	905,031		
Doctor's Offices (621111)	0.227***	0.025**	0.026**	11%	89%
	(0.026)	(0.010)	(0.010)		
	360,097	383,801	360,097		
Hospitals (622)	0.008	0.000	-0.002	-26%	126%
	(0.005)	(0.003)	(0.003)		
	381,971	757,031	381,971		
Nursing Facilities (623)	0.005	-0.002	-0.001	-16%	116%
	(0.011)	(0.006)	(0.006)		
	110,191	183,575	110,191		
Family Services (624)	0.128***	0.033***	0.034***	26%	74%
	(0.013)	(0.007)	(0.008)		
	65,926	152,957	65,926		
Childcare (62441)	0.049***	0.054***	0.050***	102%	-2%
	(0.016)	(0.011)	(0.011)		
	42,360	54,401	42,360		
Education (611)	-0.049***	-0.023***	-0.025***	50%	50%
	(0.011)	(0.008)	(0.008)		
	104,528	165,407	104,528		
Finance and Management					
Banking & Credit (522)	-0.270***	-0.027***	-0.026***	10%	90%
	(0.010)	(0.007)	(0.007)		
	976,814	988,174	976,814		
Investments (523)	0.345***	-0.026	-0.023	-7%	107%
	(0.037)	(0.018)	(0.018)		
	80,185	80,349	80,185		
Insurance (524)	-0.068	-0.038	-0.038	55%	45%
	(0.107)	(0.085)	(0.082)		
	567,213	567,270	567,213		
Administration (561)	0.096***	0.039**	0.039**	40%	60%
	(0.024)	(0.014)	(0.013)		
	2,279,915	2,280,313	2,279,915		
Utilities (221)	0.118	0.012	0.101	86%	14%
	(0.251)	(0.188)	(0.199)		

	Nonprofit Estimates	Differential		Difference Decomposition	
	Cross Sectional Difference	Within-Worker Estimate	Gen-Diff-in-Diff Estimate	Percent Demand	Percent Supply
	(1)	(2)	(3)	(5)	(6)
	161,157	161,853	161,157		
Classic Charities					
Legal Aid (54111)	-0.256***	-0.131***	-0.129***	51%	49%
	(0.036)	(0.036)	(0.036)		
	270,984	271,996	270,984		
Religious Organizations (8131)	-0.264***	-0.099***	-0.098***	37%	63%
	(0.014)	(0.013)	(0.013)		
	4,003,999	4,012,994	4,003,999		
Grantmaking Foundations (8132)	-0.123***	-0.020*	-0.017	14%	86%
	(0.012)	(0.009)	(0.009)		
	3,943,822	3,971,955	3,943,822		
Social Advocacy (8133)	-0.243***	-0.024***	-0.021*	9%	91%
	(0.012)	(0.009)	(0.009)		
	3,949,345	3,974,332	3,949,345		
Civic Organizations (8134)	-0.081***	-0.051***	-0.050***	62%	38%
	(0.015)	(0.011)	(0.011)		
	2,581,211	2,593,224	2,581,211		

Notes: Table is based on the estimation of equation (2) estimated for various industries. Column (1) reflects the cross-sectional earning difference between nonprofit and for-profit workers in each category. Column (2) estimates the nonprofit penalty by adding worker fixed effects to the estimation of (1). Column (3) implements a generalized difference in difference which adds to the specification in column (2) event-time fixed effects to account for general dynamics surrounding job changes. Column (4) reflects an estimate of how much of the nonprofit differential arises from demand-side forces, calculated by dividing the value in column (3) with the value in column (1). Column (5) reflects the share of the nonprofit differential arising from supply differences, which is the remaining nonprofit differential unexplained by demand. The third row for each estimate provides the relevant N. (* $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$)

Figure 1.1: Event Study - Job Changes from For-profit Employers

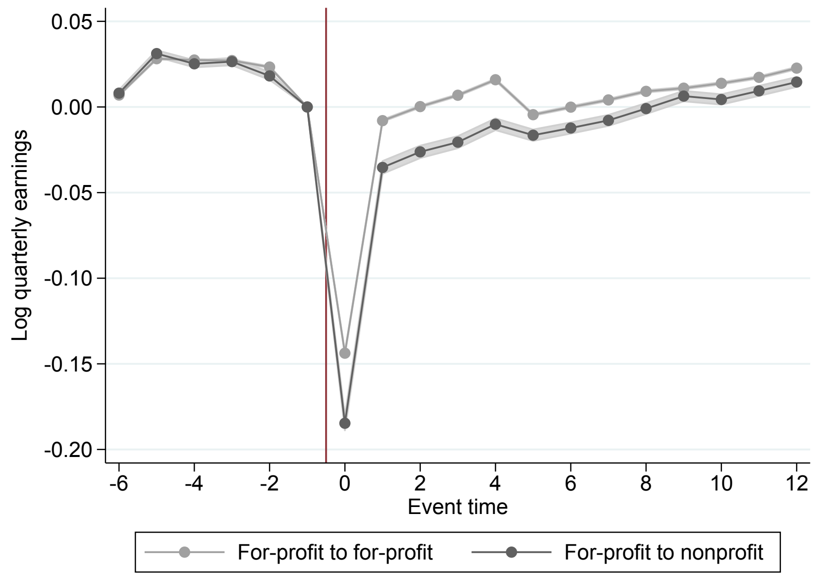

Notes: We plot the event-time dummies for workers who changed jobs between 2003-2012 and held the previous and new job for at least six quarters (eighteen months). After t=6, the results derive from an unbalanced panel. Controls include a full set of event-time dummies, year-quarter dummies, and event-specific dummies, a refinement of worker FE.

Figure 1.2: Nonprofit Influence on the Income Distribution

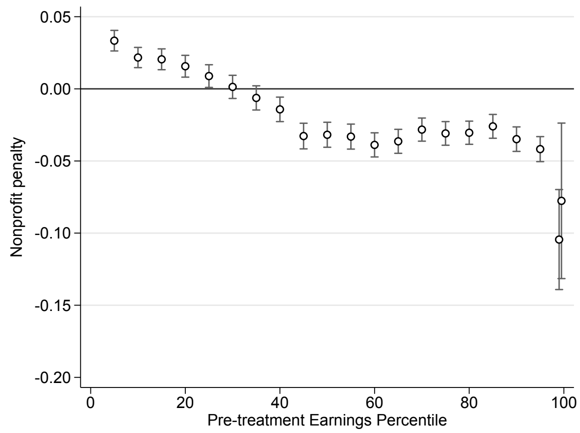

Notes: Figure shows the coefficients on the nonprofit indicator from equation 1.2 for several pre-treatment income groups. To determine pre-treatment wage groups, we residualize log quarterly earnings from the pre-treatment year on event and industry fixed effects. We use this residualized log quarterly earnings to partition workers into pre-treatment wage quantiles. The data are from administrative unemployment insurance records for the universe of Florida workers from 2003 through 2012.

Figure 1.3: Classic Nonprofit Event-Study Figures

(a) Panel A: Legal Employers

(b) Panel B: Religious Employers

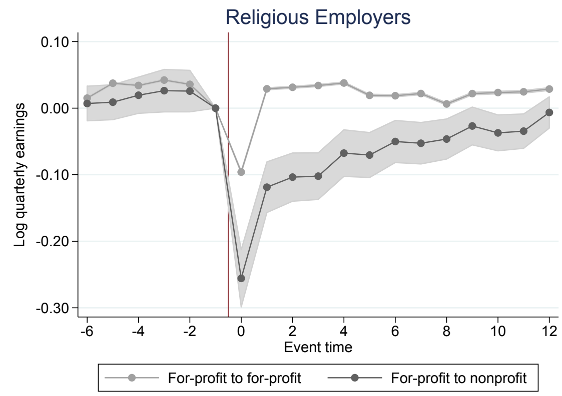

(c) Panel C: Civic Organization Employers

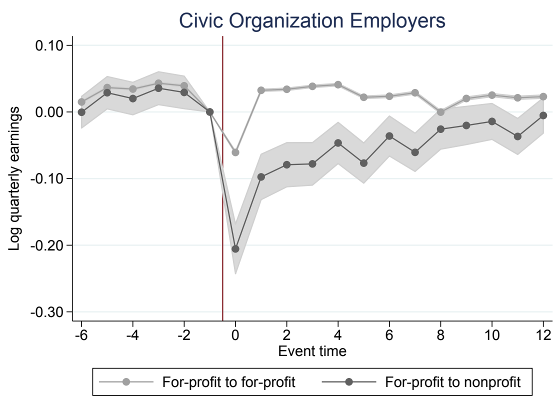

(d) Panel D: Social Advocacy Employers

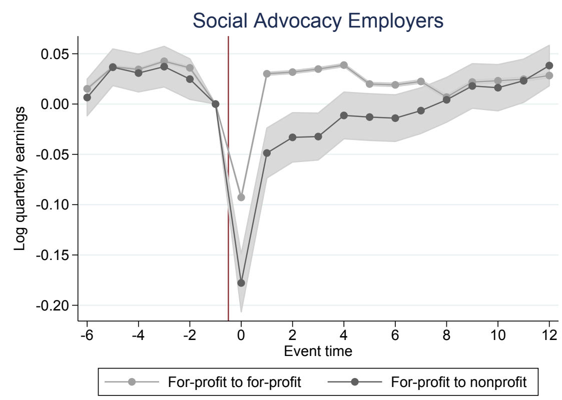

Notes: Figure shows the coefficients α_j^q in equation 1.1 for two event types: moves from a for-profit to a nonprofit (which we refer to as the treatment group) and moves from a for-profit to another for-profit firm (the control group) for various 3-digit NAICS industries. The dependent variable is log quarterly earnings. The event-time dummy at $t = -1$ is omitted. To generate the control event for religious, civic organization, and social advocacy industries, we identify the three 3-digit NAICS codes that most commonly transition to that particular nonprofit type and identify workers transitioning between for-profit jobs in those industries. The grey, shaded areas bounding each line represent the 95-percent confidence interval.

Figure 1.4: Commercial Nonprofit Event-Study Figures

(a) Panel A: Education Employers

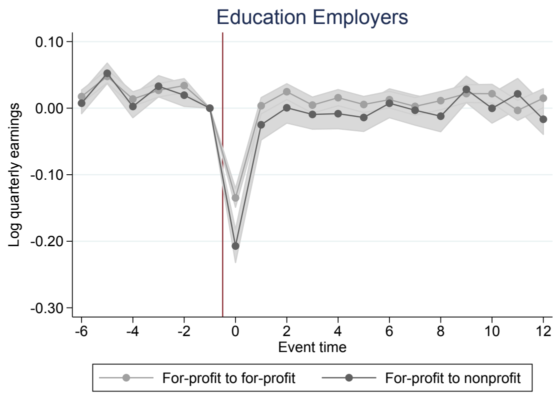

(b) Panel B: Nursing Care Employers

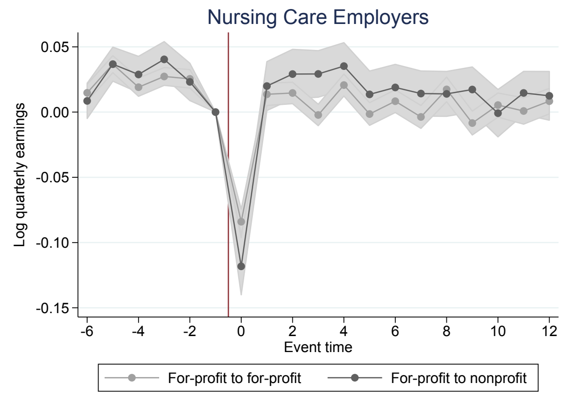

(c) Panel C: Outpatient Healthcare Employers

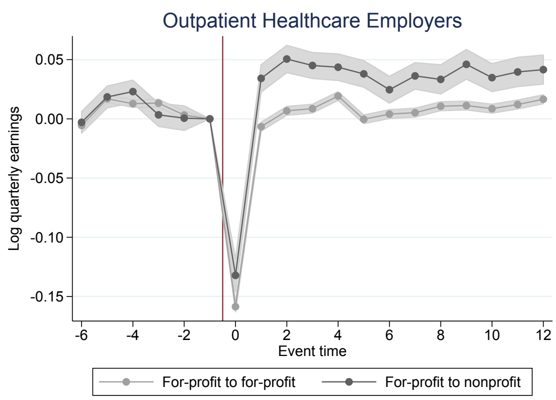

(d) Panel D: Doctor's Office Employers

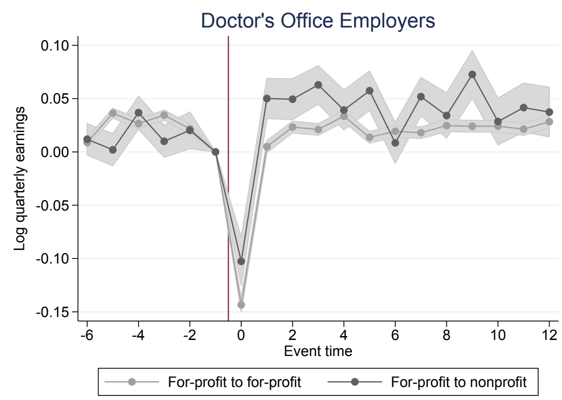

Notes: Figure shows the coefficients α_j^q in equation 1.1 for two event-types: moves from a for-profit to a nonprofit (which we refer to as the treatment group) and moves from a for-profit to another for-profit firm (the control group) for various industries. The dependent variable is log quarterly earnings. The event-time dummy at $t = -1$ is omitted. All industries are determined by 3-digit NAICS codes, except for Doctors' offices, which corresponds to a 6-digit NAICS code. The grey, shaded areas bounding each line represent the 95-percent confidence interval.

Figure 1.5: Nonprofit Differential and Nonprofit Penalty Over Time

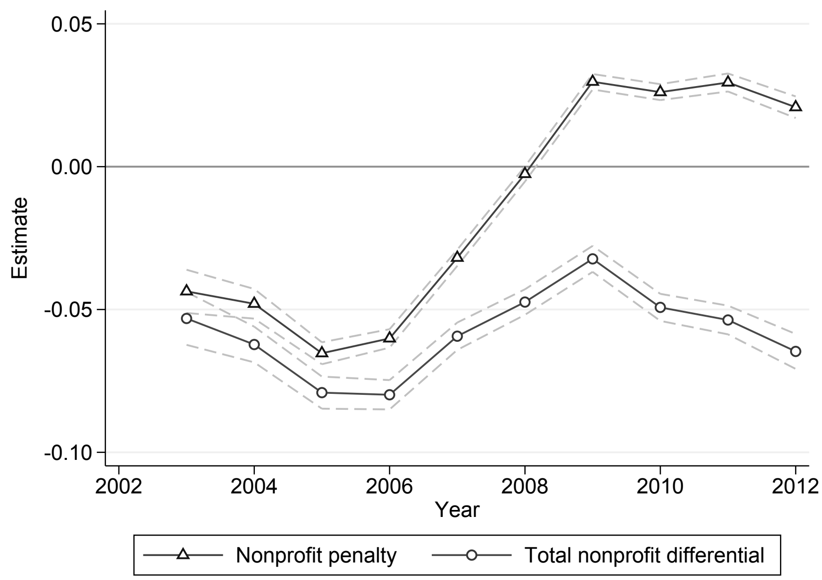

Notes: The circle-dotted line represents the cross-sectional nonprofit differential in each year. The triangle-dotted line presents the nonprofit penalty in each year, which accounts for worker-specific differences using individual fixed effects.

Figure 1.6: Work Type over the Business Cycle

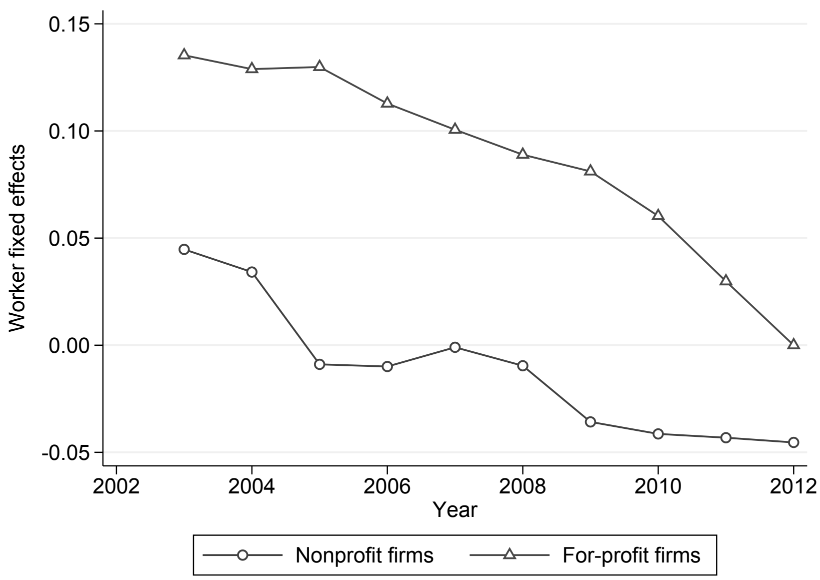

Notes: Figure shows the worker fixed effects plotted over time. Worker fixed effects are estimated from an AKM worker-firm fixed-effects model with year and quarter controls. All sample restrictions described in section III are imposed.

Chapter 2

Price and Quantity Impacts of the Earned Income Tax Credit on a Large Durable Goods Market

2.1 Introduction

The first chapter of this dissertation estimates the effect of a firm's tax-exemption status on workers' wages. Given that wage changes are infrequent, often large, and lasting, this change in wages has long-term implications for the worker. This chapter complements chapter 1 by also examining the effect of a subsidy on an economic actor. In the first chapter the intended target of the subsidy was a firm and my investigation focused on the effect of the tax-break on related economic actors, the workers. I focused my first chapter to investigating individual outcomes.

This chapter also focuses on a government subsidy, but shifts the focus to a setting with an immediate short-term effect. EITC recipient consumption patterns change quickly after the receipt of a subsidy, but the change is fleeting. In addition to exploring individual outcomes, this chapter also analyzes market changes, something chapter 1 did not do extensively. Both chapters take advantage of universal records from large US states.

2.2 Background

The EITC has been credited as one of the most successful anti-poverty social programs in the United States, lifting 9.1 million people out of poverty (Nichols and Rothstein, 2015). Transfers from the EITC are potentially large, and a low-income family with two children that earned between \$14,000 and \$22,000 in 2015 was eligible to receive over \$5,500 with their tax return, or 25-40% of their adjusted gross income.

Unlike many social programs where benefits are transferred monthly, the EITC is a tax credit that is transferred annually as a single lump-sum upon receipt of a federal tax return.

When the program was first introduced, recipients were given the option to smooth their EITC transfer throughout the year, but nearly all recipients chose to receive their benefit in a single, large lump-sum transfer. Some have suggested that this single transfer provides a savings commitment device, allowing potentially liquidity constrained recipients to purchase consumer goods and large durable goods, which would otherwise be infeasible (Nichols and Rothstein, 2015). Purchases of furniture, appliances, household goods, and vehicles have been seen to spike among potential EITC recipients in February, when the majority of recipients receive their refunds Barrow and McGranahan (2000), Goodman-Bacon and McGranahan (2008)).

Most of the previous work has examined the demand impacts of the EITC in a partial equilibrium framework, without considering the supply response and the ultimate economic incidence of this demand shock. A recent chapter summarizing the previous work regarding the EITC suggest that the economic incidence of the EITC is still "under-studied" (Nichols and Rothstein, 2015). In a general equilibrium setting, a tax or transfer policy that shifts out either the supply or demand curve will result in a new equilibrium price that shifts the realized incidence of the policy, despite the statutory incidence of the tax. The share of the tax (or transfer) burden shifted to the opposite side of the market will depend on the relative elasticities of supply and demand. By design, the EITC embodies a delayed wage subsidy for low wages. To this point, the previous research focusing on the incidence of EITC transfers has focused on changes in the labor market. However, as the previous research has shown that the EITC shifts out demand for consumer goods and durables (at least temporarily), this transfer could also be viewed as a subsidy on durable goods and other big ticket items given to EITC recipients. As such, it is equally important to understand to what degree the EITC benefits are captured by suppliers in these markets. Understanding this general equilibrium relationship can help evaluate the economic impact of the EITC as well as shed light on the economic incidence of low income transfer programs more generally.

2.3 Research Question

We propose to characterize the economic incidence of the EITC by asking a series of questions. First, to supplement the existing work, we ask whether the EITC shifts demand in the used car market, and if so, to what extent? After understanding from the data where demand increases, we next ask how do equilibrium prices adjust.

2.4 The Earned Income Tax Credit

Congress created the Earned Income Tax Credit in 1975 to help low-income families off-set payroll taxes while still incentivizing work. For an in-depth history of the EITC, we refer the readers to Ventry Jr (2000). The credit is structured as a percentage of income, the percentage of which increases with over a low income levels, plateaus over middling income levels, and then phases out over higher income levels. The number of children and the marital status of household also effect the amount of EITC each household receives, but the overall increase, plateau, decrease structure of the credit holds across these variations. Figure 2.2, taken from Goodman-Bacon and McGranahan (2008) shows the EITC schedule for selected years. For the year 2006, a single parent earning between \$11,340 and \$14,810 would have received the maximum credit amount of \$4,536, which is roughly a third of their earned income. The average EITC return for filers over the years of our analysis is is \$2211.

Historically, more than 75 percent of EITC recipients receive their EITC in the month of February or March. LaLumia (2013) plots the share of EITC returns issued in each month, using data from *Monthly Treasury Statements* for the years 1998 to 2007. We show her plot below in figure 2.3. Of refunds made to general filers, approximately 20 percent are made in each February, March, April, and May. In contrast, 55 percent of EITC recipients receive their refund in February and 22 percent receive their refund in March. We will exploit this narrow time window of EITC returns in our analysis of car sales and prices in the EITC return months.

In figure 2.4 below we plot the share of the zipcode population which receives the EITC in 2010 [1]. We focus only on Texas, the state in our analysis sample. Darker blues correspond to a larger percentage of the population receiving the EITC. On average, 24 percent of the zipcode population receives the EITC, although there is wide variation across the state of Texas. The standard deviation of EITC population share across zipcodes is 11 percent. There are large swaths of dark blue zipcodes in Southern Texas, and lighter swaths in Northern Texas, and central regions of the state exhibit significant variation in population shares among zipcode neighbors. This variation across Texas is pivotal to our difference-in-difference research design.

In figure 2.5, we plot the EITC dollars per household and see similar variation throughout Texas as in figure 2.4. The average EITC dollars per houshold across zipcodes is \$624, with a standard deviation of \$415. We plot this measure to give a sense of how many dollars are flowing into each zipcode, and presumably each local used car market, although our main analysis focus on the population share variation.

2.5 Model

Consider the market for a large durable good, such as cars, in a static setting. Assume that in this market there is a continuum of potential consumer households of size 1, where each household's latent demand is defined by

[1] All of our analyses map U.S. Postal zipcodes to ZCTA regions in order to align with Census shape files and county crosswalks. For simplicity, we refer to a ZCTA region as a zipcode, as over 80% of ZCTA regions exactly encompass a zipcode

$$d_i^*(p_i) = \begin{cases} 1 & \text{if} P \leq p_i \\ 0 & \text{else} \end{cases} \tag{2.1}$$

and where p_i is the highest price household i will pay for the good x. First, for simplicity, assume $p_i \sim u[0,1]$. Under a uniform distribution, aggregate latent demand for good x will be $D*(P) = 1-P$. However suppose that a share of households, s, that are cash constrained and cannot purchase x, even at the price that they are willing to pay p_i. Once again assume for simplicity that $s_i \sim u[0,1]$ and at least initially assume that s_i and p_i are uncorrelated. Given these cash constraints, realized demand becomes

$$d_i^*(p_i, s_i) = \begin{cases} d_i^*(p_i) & \text{if} s_i \geq s \\ 0 & \text{if} s_i \leq s \end{cases} \tag{2.2}$$

In aggregate, realized demand will then be $D(P) = (1-P)(1-s)$. Now suppose that everyone with $s_i \leq s$ is given a transfer of income that relaxes the constraint, such that $d_i = d_i^*$ for all i. Because we have assumed s_i and p_i are uncorrelated, the demand curve will uniformly shift out from D to D^*, depicted in figure 2.1. This outward shift in demand will lead to a new market price, that will depend on the elasticity of supply of good x. If x is supplied more elastically (curve S_1) then the new market price will be P_1. If the supply of x is relatively more inelastic (curve S_2) then the new market price will be P_2. In each case we can calculate the incidence of the transfer, by comparing the original price (P_0), to the new price (P_1, or P_2), and to the price that would have been observed in the original state if demand was at the new level of demand (p_{01} or p_{02}).

The original inverse demand function is $P = 1 - \frac{D}{1-s}$. In either setting, the consumer incidence will be

$$share_c = \frac{P_0 - p_{0i}}{P_i - p_{0i}} = \frac{P_0 - \left(1 - \frac{X_i}{1-s}\right)}{P_i - \left(1 - \frac{X_i}{1-s}\right)} \tag{2.3}$$

Note that the consumer share of the transfer corresponds to the blue arrows in the figure, and the producer's share will be $share_p = 1 - share_c$, and corresponds to the orange arrows. If the supply curve is more elastic than the demand curve, a larger share of the transfer will be retained by the consumer. However, if the supply curve is more inelastic, a larger share of the transfer will be passed through to the producers. As we do not know the relative elasticity of the demand and supply curves, this remains and empirical question for future work.[2]

[2]The assumption on the distribution of p_i and s_i, as well as the correlation between the two can be relaxed and the same general result will hold.

2.6 Data

Our data comes from two sources: 1) Texas DMV registration records for the years 2004-2010 and 2) Zipcode-level statistics on EITC measures, made publicly available through a collaboration between the IRS and the Tax Policy Center[3].

We constrain our sample of registration records to passenger vehicles which have recently been purchased. The Texas DMV records contain fields for vehicle mileage, sale price, sale date, registration date, vehicle type, make, and model, owner name and address, and seller name. We classify a vehicle as used if the odometer reads 10,000 or more miles upon registration. Texas law requires that owners register vehicles no more than 30 days after purchase and we restrict our sample to registrations that coincide with a purchase date that is less than 30 days prior to registration.

Using owner zipcode from the DMV registration records, we merge the zipcode-level records provided by the Tax Policy Center to our registration records dataset. These zipcode-level records include the total number of issued tax returns, the number of tax returns including an EITC, and the total number of EITC dollars issued. For our analysis, we convert our zipcodes to ZCTA (zipcode tabulation area) codes issued by the Census. The majority of ZCTA codes are simply the original zipcodes assigned by the US Postal service, but in some cases two or more zipcodes are combined under one ZCTA. Using the 2000 Census ZCTA measures of population and household income, we compute additional measures of EITC variation: the share of the population which receives the EITC and the number of EITC dollars per household. Summary statistics of our data are provided below in table 2.1.

2.7 Empirical Strategy: Difference-in-differences

This section describes our empirical strategy to recover causal estimates of the effect of EITC returns on quantities and prices in the used vehicle market. We use a difference-in-differences design, taking advantage of the regional and yearly variation in EITC population shares across Texas and our knowledge of the EITC return schedule. Because most EITC returns are issued in February and March, we can surmise that zipcodes with large populations of EITC recipients will be "treated" by the EITC return in February and March. The coefficients of interest from specification 2.4 below will measure the effect of EITC returns on a zipcode's used car sales.

$$
\begin{aligned}
ln(\text{sales}_{z,t}) = \beta_0 &+ \sum_{m \in \mathbf{M}} \left(\beta_m * 1\{month = m\} * \text{shareEITC}_{g,yr(t)}\right) \\
&+ \beta_2 * \text{shareEITC}_{z,yr(t)} \\
&+ \Gamma_t + \Gamma_{m(t)} + \Gamma_{yr(t)} + \Gamma_z + \epsilon_{z,t}
\end{aligned}
\tag{2.4}
$$

[3]The publicly available data can be found at https://tpc-eitc-tool.urban.org/

In this specification, $ln(\text{sales}_{z,t})$ is the log of used cars sales in zipcode z in month-year t. The variable $\text{shareEITC}_{z,yr(t)}$ is the share of the population claiming the EITC for zipcode z in year $yr(t)$. This specification includes time-period (Γ_t), month ($\Gamma_{m(t)}$), and year fixed effects ($\Gamma_{yr(t)}$), as well as zipcode fixed effects (Γ_z). The main interaction terms include month-share interactions for every month except July, which we use as our baseline. The set of months excluding July form the set of months $\mathbf{M}$. The coefficients of interest are the parameters estimated from the interaction of the EITC population share with the February and March month dummies, $\beta_{m=Feb}$ and $\beta_{m=Mar}$. These coefficients can be interpreted as a lower bound on the shift in the demand curve from our model in section 2, under the stringent assumption that the supply curve remains unchanged.[4]

Specification 2.5 below presents our empirical design to recover the changes in prices caused by EITC returns. Our design is similar to the specification 2.4, with individual log sales price as the outcome variable, instead of zipcode-level sales numbers.

$$\begin{aligned} ln(\text{price}_i) = \beta_0 + &\sum_{m \in \mathbf{M}} \left(\beta_m * 1\{month = m\} * \text{shareEITC}_{g,yr(t)}\right) \\ &+ \beta_2 * \text{shareEITC}_{z,yr(t)} \\ &+ \Gamma_t + \Gamma_{m(t)} + \Gamma_{yr(t)} + \Gamma_z + \Gamma_c + \epsilon_i \end{aligned} \tag{2.5}$$

We use the individual record level data from the Texas DMV to obtain vehicle prices and characteristics and zipcode-level EITC variation. As above, the coefficients of interest are the $\beta_{m=Feb}$ and $\beta_{m=Mar}$. We include additional controls for used car characteristics, denoted in Γ_c. These controls include vehicle model and model year interaction terms, as well as quadratic controls for odometer mileage.

2.8 Results

Figure 2.6 and 2.8 present the main results from specifications 2.4 and 2.5, respetively. These figures plots the coefficients from the interaction effects of month and EITC population shares (measured in percentages). The coefficients in figure 2.6 show the percentage change in the difference between a given month's sale and July sales, the change being driven by a 1 percentage point increase of EITC population share.

The percentage change in sales jumps from nearly 0 to 1.3 from January to February and climbs to 1.4 percent in March. It then steadily falls until June, where it hovers near zero for the rest of the year. Since most EITC returns are issued in February and March (see figure 2.3 this increase in demand in February and March confirms our hypothesis that EITC returns drive consumption in the weeks after receipt.

[4]It is a lower bound because our model assumes a lump sum endowment lifts all cash-constrained individuals onto their true demand curve, while the actual EITC only unconstrains some previously cash-constrained individuals.

Multiplying this EITC effect by the average EITC population share (24 percent) for the month of March, results in an estimated 33.3 percent increase in demand. Similarly, multiplying the effect by the 25th and 75th percentile for EITC population shares leads to percentage increases of 23.0 percent and 41.5 percent, respectively. Given the variation depicted in figure 2.4, these results imply demand for used cars is varying greatly across the state of Texas in February and March.

To get a general sense of what used car sales are doing throughout the year, we plot $\Gamma_m + s_p \cdot (\beta_m + \beta_2)$, for each month of the year. The coefficient Γ_m is the month m's fixed effect from specification 2.4, β_m is the coefficient from the interaction term $1\{month = m\} * \text{shareEITC}_{g,yr(t)}$ and β_2 is the coefficient from the EITC population share. The term s_p is the pth percentile of EITC population share. Although this exercise does not give exact sales numbers throughout the year, it does allow us to see the overall sales trends in high and low EITC population areas. In both types of areas, sales are lower in the first half of the year and higher in the second half. Reflective of results in figure 2.6, the difference in car sales between high and low EITC population areas widens in February and March and narrows to zero by June.

Similarly to quantity, prices also increase in areas with a larger EITC population shares in month February and March. Figure 2.8 plots the coefficients from the interaction effects of month and EITC population shares from specification 2.5. These coefficients show the percentage change increase in prices caused by a 1 percentage point increase in a zipcode's EITC population share. Interestingly, the percentage change in the price differences jumps in January to .08 percent and climbs slightly in February and March, reaching a peak of .12 percent for every 1 percentage point increase. The percentage change decreases steadily until bottoming out in June at -.05 percent. Changes track zero for the rest of the year. Prices are estimated less precisely than quantities, but even so, for the first four months of the year, changes are significantly different than zero at an α level of .05. Unlike sales, which jump in February, this price jump precurses the heaviest EITC return months. One possible explanation for this puzzling January price jump could be that dealerships know EITC populations are expecting large lump sums of cash in the coming months and increase prices before returns are issued. Meanwhile, EITC populations' demand is driven by the actual receipt of the EITC.

We calculate the percentage increase in prices for a zipcode at the 25th percentile, 75th percentile, and mean of the distribution of EITC population shares. For a zipcode with the 25th percentile share, or 16.5 percent, the change in prices from March to July is 2.0 percent higher than a zipcode with no EITC recipients. A zipcode with the mean share, 24 percent, experiences a price increase of 2.8 percent and a ZCTA with the 75th percentile share or 30 percent, experiences a price change of 3.6 percent. The price response is much more muted than the sales response.

These small but significant price differences are possibly underestimated. We believe true prices for some vehicle sales are not correctly reported. There is anecdotal evidence that buyers and sellers agree to under-report the price of a vehicle so the buyer does not need to pay as much sales tax upon registration [5]. The seller then receives a percentage of the saved sales tax as reward for reporting a lower price. One imagine that sellers "rule-of-thumb" for under-reporting has muted variation compared to the actual prices of cars, which would lead to an underestimated effect and also wider confidence intervals.

Relating these quantity and price results back to the model represented in figure 2.1, we conclude that quantity shifts out greatly, increasing by 33.3 percent in March in areas with an average share of EITC recipients. Given the initial results for price increases, we hesitantly conclude the supply curve of used vehicles is quite elastic, although the pre-EITC return increase of prices in January is puzzling. This could indicate seller price manipulation that we have not accounted for. Additionally, the elasticity of the supply curves could vary across zipcode's, even those with similar shares of EITC populations. Future analysis will account for dealership concentration in zipcode's when estimating price effects.

2.9 Conclusion

This chapter uses individual vehicle registration records from the Texas DMV and publicly available zipcode-level EITC measures to determine the market effects of EITC returns for used car sales. The findings indicate that demand for used cars increases greatly in EITC heavy areas during the months of February and March while prices only increase slightly.

This chapter presents evidence that the lump-sum distribution method of the EITC achieves policymakers' goals of providing low-income households with a windfall without having an onerous amount of the transfer passed to suppliers of durable goods. Further investigation of the elasticity of the supply curve and the slight increase in January prices is warranted.

The increases in demand reflect many results of previous work regarding durable good consumption and EITC returns (see (Goodman-Bacon and McGranahan, 2008) and (Barrow and McGranahan, 2000)). To the best of our knowledge, this chapter is the first to examine the effect of EITC returns on prices as well as on quantities.

[5] Several online internet forums and newspaper advice columns cover this exact situation. An especially trafficked instance is `https://www.reddit.com/r/Autos/comments/1bjkej/buyer_wants_to_underreport_cost_of_car/`

2.10 Tables and Figures

Table 2.1: Summary Statistics

Variable	Mean	Std. dev	Min	Max
Total tax returns	5332.68	6578.57	10	91728
Total EITC returns	1264.95	1841.24	0	18499
Total EITC dollars	3016712.25	4804221.93	0	57103753.53
Average EITC pop share	23.76	11	0	75.55
Average EITC return	2211.19	319.33	0	3841.56
Average EITC $ Per Household	624.37	414.69	0	3092.5
Average household income	46366.66	19453.18	0	243562
Number of zipcodes	1, 180	1, 810	1, 810	1, 810

Note: Unweighted averages and standard deviations of zipcode-level EITC variables. Minimum and maximum values are also reported. All dollar measures are adjusted to 2010 dollars.

Figure 2.1: Two Different Supply Elasticities

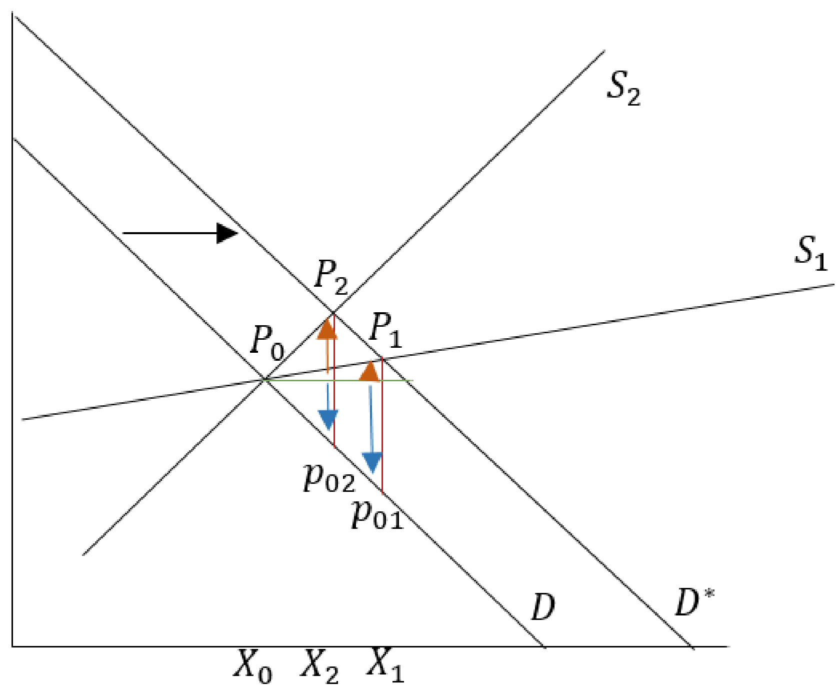

Notes: The curve $S1$ represents an elastic supply curve of used cars. The curve $S2$ represents a more inelastic supply curve. Curves D and D^* represent demand before and after the transfer, respectively. For both curves, the consumer share of the transfer corresponds to the blue arrows. The producer's share corresponds to the orange arrows. If the supply curve if more elastic than the demand curve, as in the case with $S1$, a larger share of the transfer will be retained by the consumer.

Figure 2.2: EITC Schedules

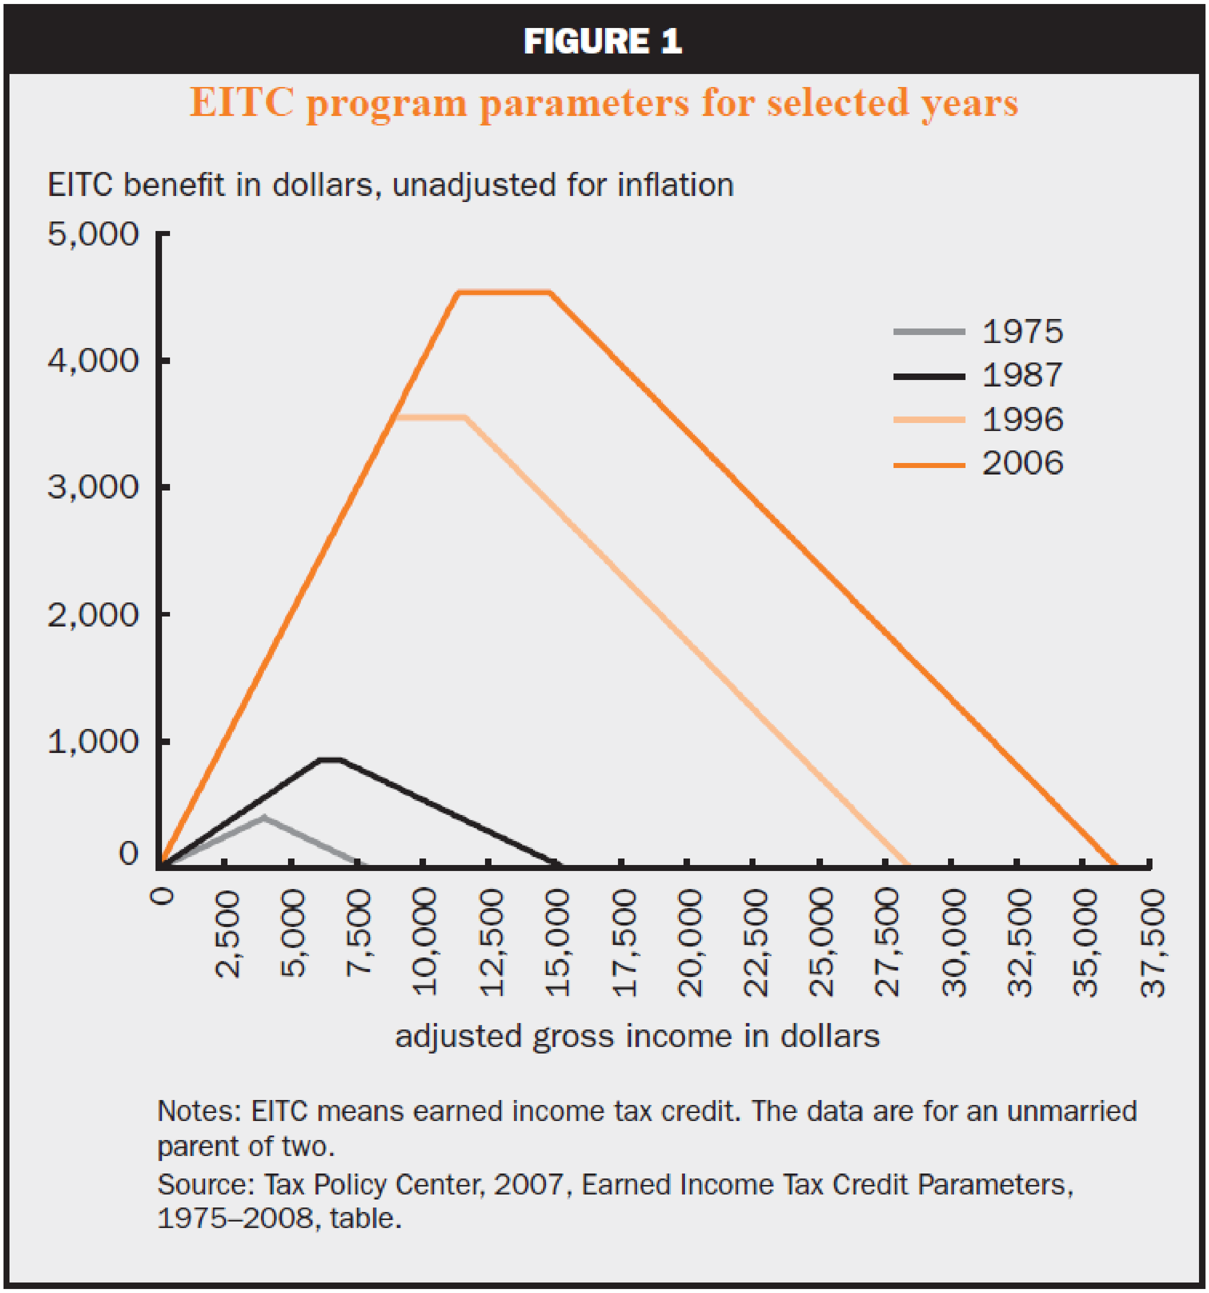

Notes: This figure is taken from Goodman-Bacon and McGranahan (2008).

Figure 2.3: Tax Refunds by Month

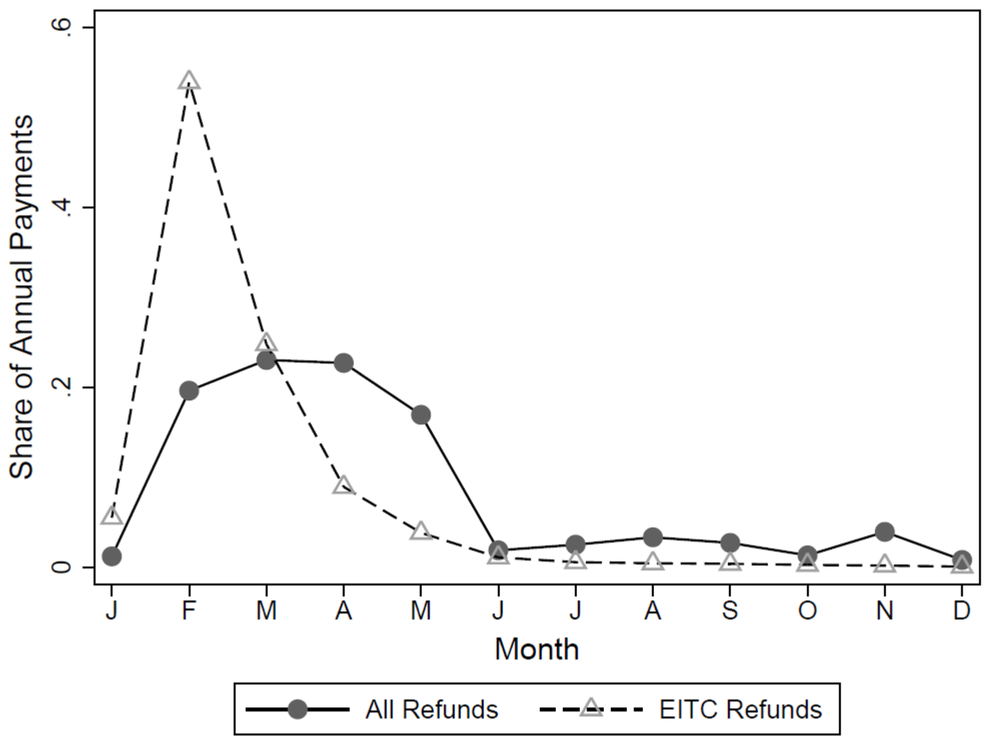

Notes: This figure is taken from LaLumia (2013)

Figure 2.4: Zipcode Share of Population Receiving EITC

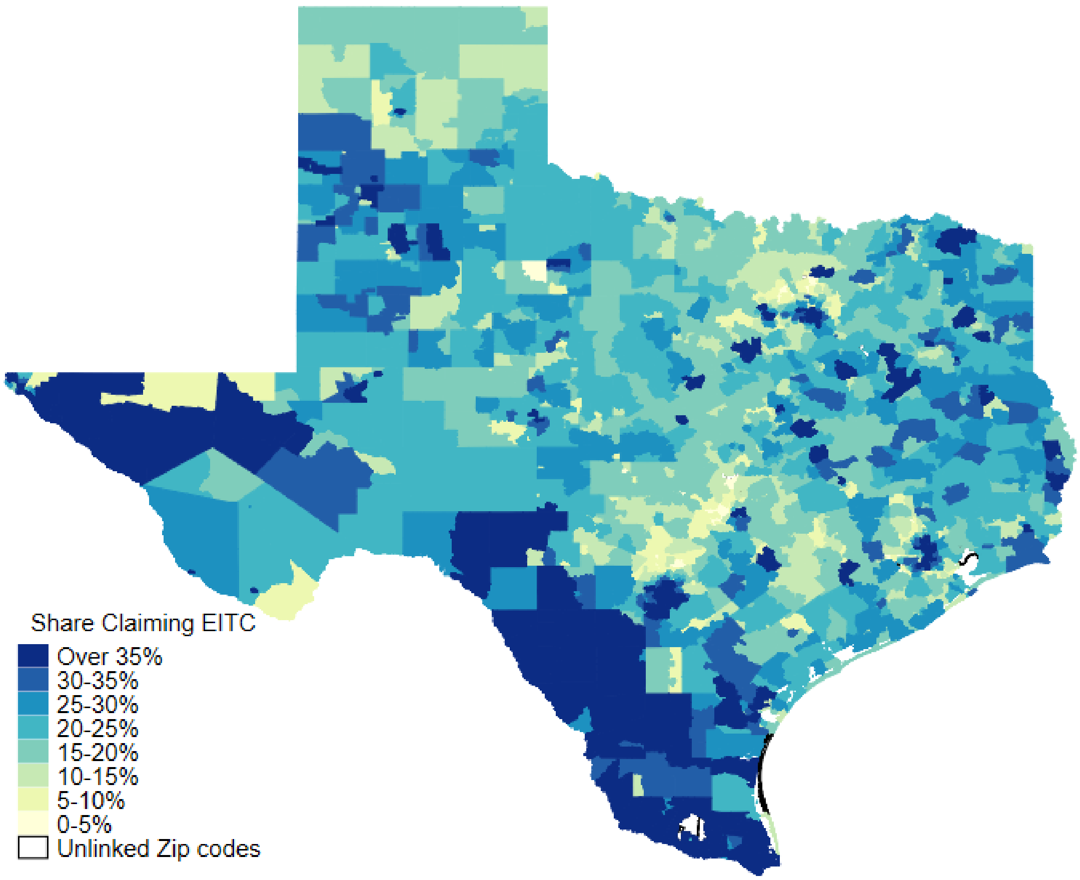

Notes: EITC zipcode level data for the year 2010 is mapped to Census ZCTA shape files. Darker regions indicate higher concentrations of EITC receiving populations.

Figure 2.5: Average EITC Dollars per Houshold, Zipcode Level

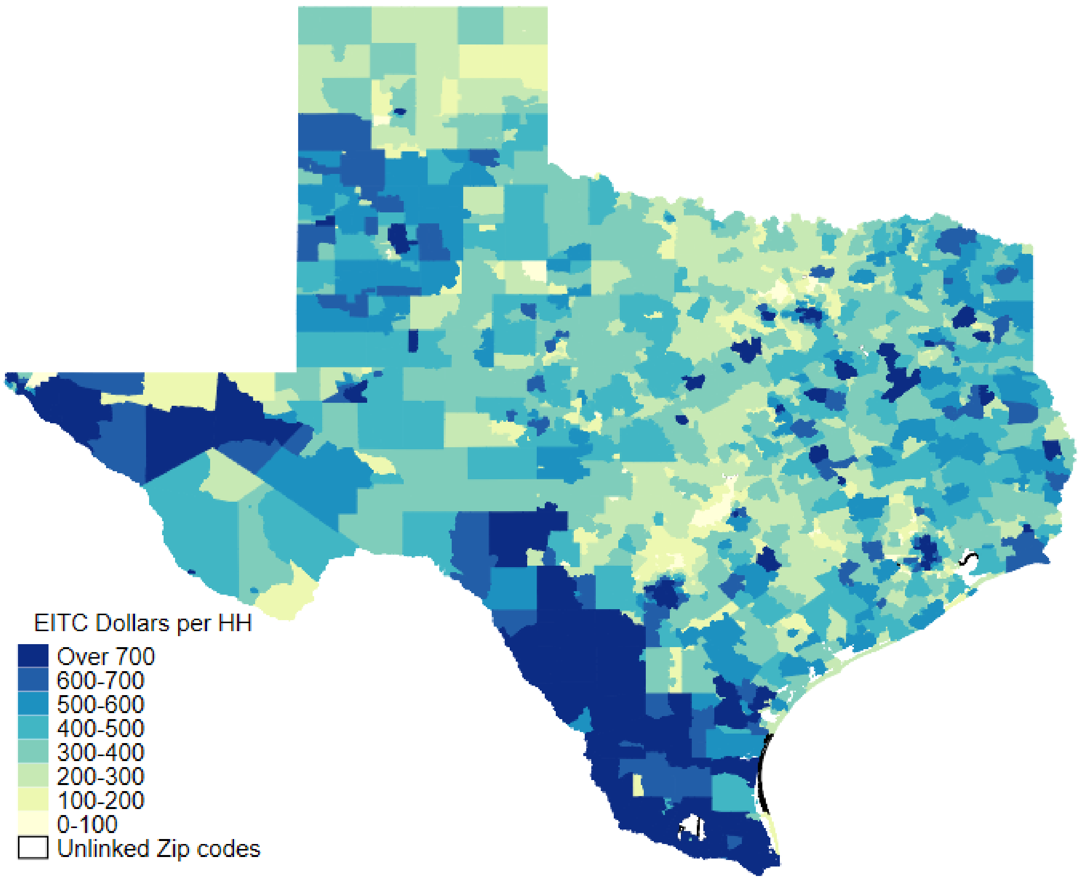

Notes: EITC zipcode level data for the year 2010 is mapped to Census ZCTA shape files. The number of households for each zipcode is obtained from the 2010 Census. Note the number of households includes all households in the zipcode, not just EITC receiving households. Darker regions indicate a higher average EITC dollars per household.

Figure 2.6: Effect of EITC Population Share on Log Sales

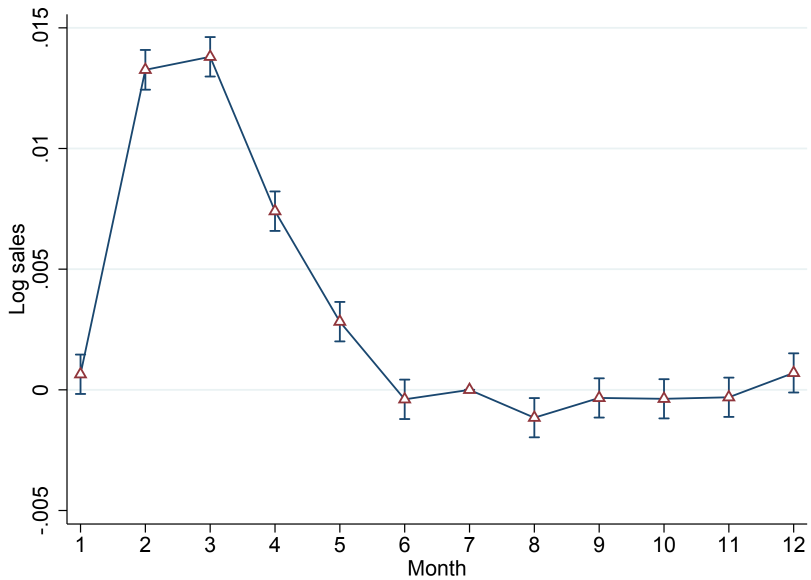

Notes: Coefficients from the month and EITC population share interaction terms from specification 2.4. Regression includes controls for month-year and zipcode. July is the baseline month. The 95 percent confidence intervals are plotted for each month.

Figure 2.7: Log Sales Trends

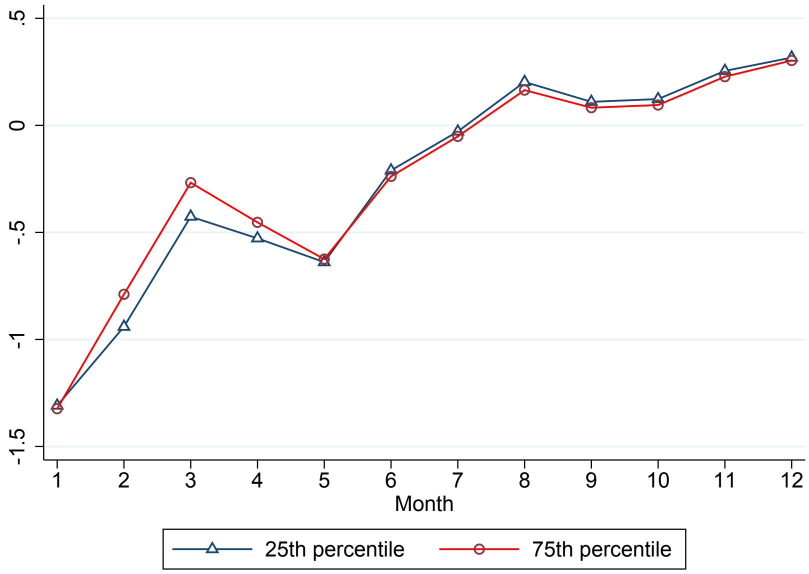

Notes: Each month's parameter is defined as $\Gamma_m + s_p \cdot (\beta_m + \beta_2)$. The coefficient Γ_m is the month m's fixed effect from specification 2.4, β_m is the coefficient from the interaction term $1\{month = m\}*$shareEITC$_{g,yr(t)}$ and β_2 is the coefficient from the EITC population share. The term s_p is the pth percentile of EITC population share. In this plot s_p takes on two values, the 25th percentile of the EITC population share (16.5 percent) and the 75th perecentile of the EITC population share (29.6 percent).

Figure 2.8: Effect of EITC Population Share on Log Prices

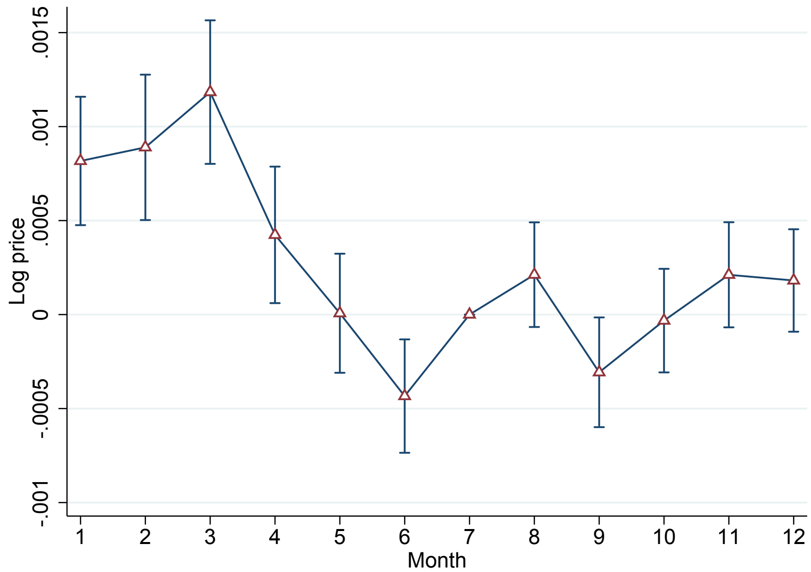

Notes: Coefficients from the month and EITC population share interaction terms from specification 2.5. Regression includes controls for month-year, vehicle model and model year, mileage, and zipcode. July is the baseline month. The 95 percent confidence intervals are plotted for each month.

Chapter 3

A Generalized Ordered Probit Model

3.1 Introduction

When studying questions related to government subsidies, publicly available data is often binned to preserve the privacy of citizens while still offering useful data to the researcher. Partial effects estimation using ordered outcomes must be treated with care to avoid misspecified and biased estimators. The first two chapters studied empirical questions regarding subsidies. This chapter complements the previous two by adding an improved, tractable ordered response model to the applied labor economist's toolkit.

Examples of datasets including binned outcomes include the Earned Income Tax dataset released yearly by the IRS. This dataset is used in chapter 2 of this dissertation to estimate the effect of the EITC on recipient consumption patterns. Census data also includes binned outcomes to protect the identity of individuals living in geographic regions. These data are used by economists in large numbers.

3.2 Background

Ordered choices are usually modeled with an ordered probit model which assumes the errors are distributed normally, or an ordered logit model which assumes the errors to be distributed logistically. These distributional assumptions impose structure which allows one to recover partial effects. Although these assumptions facilitate estimation, assuming normally or logistically distributed errors may not capture important aspects of the error distribution. For example, if one assumes errors to be normally distributed then the true distribution is assumed to be completely described by two parameters, the mean and variance and fails to account for possible skewness or thick tails that are inconsistent with the assumption of normality.

Econometric researchers have used semi-parametric ordered response estimators to avoid imposing strict assumptions on the distributional form of the errors. However, these semi-parametric estimators are often reserved for large sample analysis due to their efficiency properties. Additionally, kernel estimators require the researcher to impose some type of structure, such as specifying a kernel density and the window size, or the degree of the approximating distributional polynomial. Our approach is to assume a more flexible parametric assumption for the distribution of error terms. Similar to kernel estimators, this approach allows one to glean more information from the data regarding the unknown distribution of the error terms. Additionally, this approach could yield estimators with better properties than kernel estimators if the assumed flexible parametric distribution can accurately approximate the true distribution.

Researchers have suggested a variety of alternatives to the ordered probit and logit models. Though non-parametric ordered response models appear in theoretical econometric research (Matzkin, 1993), these models have not been widely used in empirical work (Greene and Hensher, 2010). Klein and Sherman (2002) develop one of the earliest consistent semi-parametric estimator for an ordered response model. Coppejans (2007) derives an efficiency bound for an ordered response model estimator in which the error term is unknown and develops a semi-parametric estimator that is efficient under some slightly less strict conditions than those in Klein and Sherman (2002). Chen and Khan (2003) focus on cases in which the error distribution is known, but heteroskedasticity is present. Lewbel et al. (2003) adds to the work of Chen and Khan (2003) by considering a case in which the error distribution is unknown, and developing a semi-parametric estimation framework. Lewbel et al. (2003) both develops a threshold estimation that is consistent in a generalized setting, and shows that binary choice location estimators can be converted into threshold estimators. Other chapters that outline approaches to increasing the flexibility of ordered response models by incorporating non-parametric covariate effects in an additive regression framework include Yee and Wild (1996), Donat and Marra (2017), Wood, Pya and Säfken (2016). Rigby and Stasinopoulos (2005) and Donat and Marra (2018) explore some alternative distributions. Additionally, Mora and Moro-Egido (2008) outline some specification tests to test for non-linearity, heteroskedasticity, and the validity of the assumed error distribution. This chapter contributes to the literature by presenting a simple, tractable, yet powerful option to estimating ordered response models using flexible distributions from the skewed generalized t family Theodossiou (1998).

It is important to explore improvements to the ordered probit and ordered logit models because results from ordered choice estimation can have large policy effects. Some examples include Pudney and Shields (2000) who study lifetime earnings differ- ences between ethnic groups. Lemp, Kockelman and Unnikrishnan (2011) analyze the effect of the severity of car crashes on injuries received, Chiswick (1991) studies the determinants of English proficiency among low-skilled immigrants, and Zweimüller and Winter-Ebmer (1994) study the gender wage gap in the private and public sectors. Determinants of life or job satisfaction are also often estimated using ordered choice models. An empirical application modeling the relationship between life satisfaction, income and demo- graphic variables is included in this chapter and is modeled after the application in Hodge and Shankar (2014).

The rest of the chapter proceeds as follows: Section 3.3 presents the model, Section 3.4 discusses an empirical application and corresponding results, Section 3.5 compares alternative estimators using Monte Carlo simulations and Section 3.6 concludes.

3.3 The Model

Consider the model

$$y_i^* = x_i\beta + \epsilon_i \qquad 1 \leq i \leq N \tag{3.1}$$

where y_i^* is a continuous, latent variable, x_i is a 1 x k vector of explanatory variables, is a k x 1 vector of unknown parameters and the ϵ_i are assumed to be independently and identically distributed with a pdf denoted $f(\epsilon, \theta)$ with distributional parameters, θ. We could assume heteroskedasticity, as in Chen and Khan (2003); however, to focus on a comparison of the benefits of using flexible error distributions, we will restrict the discussion in this chapter to the homoskedastic case. The number of observations is N. The above model could be consistently estimated using OLS if y_i^* was observed and $E[\epsilon_i|x_i] = 0$. Often this is not the case.

In ordered response models, the true, continuous variable y_i^* is represented by the variable y_i, which only spans a finite number of totally ordered outcomes. Due to the discrete-choice nature of the data, OLS will result in heteroskedastic errors and predicted probabilities that may fall outside the range of (0, 1) for each outcome described in the vector y. To circumvent this problem, maximum likelihood estimation is often used to estimate the unknown parameters. Consider the observed variable y_i:

$$y_i = \begin{cases} 1 & \text{if } y_i^* < \alpha_1 \\ 2 & \text{if } \alpha_1 \leq y_i^* < \alpha_2 \\ 3 & \text{if } \alpha_2 \leq y_i^* < \alpha_3 \\ \vdots & \\ J & \text{if } \alpha_{J-1} \leq y_i^* \end{cases} \tag{3.2}$$

where J is the number of mutually exclusive categories of y_i. The probability of observing a particular outcome, for $1 \leq i \leq J$ is given by

$$\begin{aligned} Pr(y_i = j|x_i) &= Pr(\alpha_{j-1} \leq y_i^* < \alpha_j) \\ &= Pr(\alpha_{j-1} - x_i\beta \leq \epsilon_i < \alpha_j - x_i\beta) \\ &= F(\alpha_j - x_i\beta; \theta) - F(\alpha_{j-1} - x_i\beta; \theta) \end{aligned} \tag{3.3}$$

where F is the cumulative distribution function for ϵ_i, $\alpha_0 = \alpha_{1-1} = -\inf$ and $\alpha_j = \inf$. The presence of F leads to a maximum likelihood estimation framework. If we define

$$z_{ij} = \begin{cases} 1 & \text{if } y_i = j \\ 0 & \text{else} \end{cases} \tag{3.4}$$

then we can write the log-likelihood function as follows:

$$logL = \sum_{i=1}^{N}\sum_{j=1}^{J} z_{ij} log\left[F(\alpha_j - x_i\beta;\theta) - F(\alpha_{j-1} - x_i\beta;\theta)\right] \tag{3.5}$$

This log likelihood function is maximized with respect to β, θ, and the cut points $\alpha_1 < \ldots < \alpha_{J-1}$. In the case of two discrete outcomes, the loglikelihood function in 3.5 simplifies to the binary choice model with one cut point which is normally set to be 0 to achieve identification of the intercept term.

The most common choices for the cumulative distribution function of the errors, F, are the cumulative normal and cumulative logistic distributions. Although a maximum likelihood model framework which assumes errors are either normal or logistically distributed provides for improvements relative to OLS, misspecification of the error distribution can lead to inconsistent and biased estimators (Greene and Hensher, 2010; Klein and Sherman, 2002).

Allowing the errors to take on a more generalized distributional specification has potential to reduce the problem of distributional misspecification and lead to more efficient estimators. The Skewed Generalized T distribution (SGT) is an example of a more generalized distribution. While there are other distributions, such as a mixture of normals, that can account for diverse distributional characteristics such as asymmetry or thick-tails, we use the SGT distribution because it nests many commonly used distributions and can accommodate a wide range of skewness and kurtosis (Kerman and McDonald, 2013). Figure 3.1 illustrates the relationship between the SGT and many of its special and limiting cases corresponding to different parameter values.

The SGT was introduced by Theodossiou (1998). Its cumulative distribution is given by

$$SGT(\epsilon, m, \lambda, \sigma, p, q) = \frac{1-\lambda}{2} + \frac{(1+\lambda sign(\epsilon - m))}{2} sign(\epsilon - m) B_z(1/p, q) \tag{3.6}$$

where the incomplete beta function is represented by B_z and z is given by

$$z = \frac{|\epsilon - m|^p}{|\epsilon - m|^p + q\sigma^p(1+\lambda sign(\epsilon - m))^p} \tag{3.7}$$

The parameter m is a location parameter, λ controls skewness, and p and q are positive shape parameters that determine peakedness and kurtosis. Letting the parameter q in the SGT grow indefinitely large results in the skewed generalized error distribution (SGED). The cumulative distribution for the SGED is given below.

$$SGED(\epsilon, m, \lambda, \alpha, p) = \frac{1-\lambda}{2} + \frac{1+\lambda sign(\epsilon - m)}{2} sign(\epsilon - m)\Gamma_z(1/p) \tag{3.8}$$

where Γ_z is the incomplete gamma function and

$$z = \frac{|\epsilon - m|^p}{\alpha^p(1+\lambda sign(\epsilon - m))^p} \tag{3.9}$$

The distributions are symmetric if $\lambda = 0$. For example, when $\lambda = 0$ the SGT and the SGED yield the generalized t (GT) introduced by McDonald and Newey (1988) and generalized error distribution (GED), respectively. Letting $p = 2$ in the SGT yields the skewed t (ST) Hansen (1994). To obtain the skewed Laplace (SLaplace), set $p = 1$ in the SGED, and to obtain the skewed Normal (SNormal) set $p = 2$. The distributions featured in this chapter are the SGT, SGED, SNormal, GED, SLaplace, Laplace, and Normal. These distributions were selected after a careful investigation of several of the distributions in the SGT distribution tree. They are among the more flexible distributions, and they yield the most interesting results as they allow us to compare the advantages of relaxing different parameterization assumptions. Our empirical exercise will show that using these flexible distributions yields significantly different results.

3.4 Empirical Application

Using data from the World Values Survey, waves 1–5 (the waves begin in 1981 and end in 2008) we examine the impact of income and religious activity on life satisfaction. Every wave can be thought of as a cross section and we are using all available cross sections up to this point. Although the waves span 5 different years, this is not a panel data set.

Table 3.1 has summary statistics of the dependent and independent variables used in this chapter's analysis. This application is very similar to the application used in Hodge and Shankar (2014). The dependent variable is life satisfaction (Sat), as reported by the participants. A rating of 0 indicates a dissatisfied attitude towards life and a rating of 9 indicates a very satisfied outlook on life. The variable religious is a binary variable equal to 1 if the individual considers themselves a practitioner of an established religion. Similarly, the variable high school is equal to one if the respondent has completed secondary school. Male, married and unemployed are also binary variables. Thus, the mean of the binary variables in Table 3.1 estimate the fraction of the population indicating they are religious, male, married, unemployed, and has completed secondary school. Income is a categorical variable partitioning income into ten levels. Each country's administrators of the World Values Survey decide the cutoffs for the ten income levels each year attempting to split the country into ten equally populated deciles of income. Though this formulation may involve various statistical issues, they will not be addressed in this chapter. The model we wish to estimate is

$$\begin{aligned} Sat_i^* = {} & \beta_1 religious_i + \beta_2 male_i + \beta_3 income_i + \beta_4 age_i + \beta_5 age_i^2 \\ & + \beta_6 married_i + \beta_7 unemployed_i + \beta_8 highschool_i + \epsilon_i \end{aligned} \tag{3.10}$$

with the objective of the chapter being to explore the impact of distributional assumptions on parameter estimates. We assume different distributions of the error term ei including the SGT, SGED, SLaplace, Skewed Normal, GED, Laplace, Logit, and Normal. Table 3.2 reports the estimated coefficients for the individual variables, the corresponding standard errors, the estimated distributional parameters and the log-likelihood values along with values for the corresponding Bayesian Information Criterion (BIC) for all of the distributions except for the Skewed Normal and GED. Each of the estimated pdfs are standardized with a mean of zero and unitary variance with the values (fixed or estimated) for the other distributional parameters being given at the bottom of Tables 3.2. Table 3.3 reports likelihood ratio tests between the different models.

The estimations were performed using a program written in Python [1]. The optimization procedure first used Nelder-Mead search algorithm and switched to the Powell algorithm if convergence had not been achieved after 20,000 function iterations. The Hessian matrices to calculate the standard errors were evaluated analytically. Optimized values for the simpler specifications (lower on the tree in Figure 3.1) were used as initial values to facilitate estimating the more flexible specifications. If no simpler specification existed for a distribution, optimized values were taken from an OLS regression. Because the focus of the chapter is on estimating the coefficients, the cut points are not reported in Table 3.2, but for every estimation they were monotonic.

These results show that satisfaction increases with religiosity, higher income, and being married at the .01 level of significance in each estimation. Furthermore, satisfaction is lower for males at the .01 level of significance in each estimation. The statistical significance of education and employment status depends on the error distribution with high school education being statistically insignificant with the probit and significant at the .01 level on all other estimations in Table 3.2. Unemployment is statistically insignificant in the probit and SGED estimations, significant at the.1 level with the SLaplace estimation, significant at the .05 level with the SGT estimation and significant at the .01 level on the Laplace estimation. The results in Table 3.2 suggest that satisfaction and age have a parabolic relationship, with satisfaction decreasing in age until approximately age 50 and then increasing.

In the results given above, the log likelihood values are seen to increase significantly as distributions become more flexible. The likelihood ratio tests between the SGT and less flexible distributions are given in Table 3.3. The likelihood ratio test is a statistical test used to compare the goodness of fit of an unrestricted model with one of its special or limiting cases. If l and l^* denote the log likelihood values of the unrestricted and restricted models, then the likelihood ratio test statistic, defined by $LR = 2(l - l^*)$ has an asymptotic Chi-square distribution with r degrees of freedom where r denotes the number of independent restrictions. Based on the likelihood ratio test, we can reject the hypothesis that the SGT is observationally equivalent to the Normal, Laplace, Slaplace and SGED distributions at the .0001 level. Hence it appears that the additional flexibility of the SGT distribution may be beneficial. An alternative measure that is commonly used when comparing nonnested models is the Bayesian Information Criterion ($BIC = kln(n) - 2l$) which takes account of sample

[1] The publically available scripts used for estimation include programs that do not require parallel processing and is available at https://bitbucket.org/cjohnst5/generalizedorderedprobit/src/master/. Under the restriction that the code do not run in parallel, the authors are still working on optimizations for the SGT. Estimation in both programs imposes monotonicity of the cut points

size (n), goodness of fit (l) and includes a penalty for the number of estimated parameters (k). The BIC provides additional support for the SGT.

The results in Table 3.2 demonstrate whether or not a given variable is statistically significant, but simply viewing the estimated coefficients $\widehat{\beta}$ provides less information than in a linear regression model about the marginal effect of the independent variables on the dependent variable. To fully understand the differences between two ordered response estimates, we consider the estimates of $\widehat{\beta}$, $\widehat{\alpha}$, and $\widehat{\theta}$ and explore the marginal effects.

In a linear regression model, the β vector estimates the marginal impact of changes in the independent variable on the expected value of y. However, in an ordered response model, the outcome variable is not a continuous variable, but rather an ordered set of discrete outcomes. Thus, the expected value of changing the independent variable on the outcome variables can be thought of as the expected value of changing the independent variable on the probability of each discrete outcome occurring. We find this marginal effect by taking the derivative of Equation 3.3. Notice that the marginal effect of changing x on the probability of some outcome y_i is as follows:

$$\frac{\delta P(y=j|x)}{\delta x} = \begin{cases} -\beta f(\alpha_1 - x\beta) & \text{if } j=1 \\ \beta(f(\alpha_{j-1} - x\beta) - f(\alpha_i - x\beta)) & \text{if } 1<j<J \\ \beta f(\alpha_{J-1} - x\beta) & \text{if } j=J \end{cases} \tag{3.11}$$

where $f()$ is the pdf of the error distribution. Unlike the error distribution in linear regression models, the error distribution in the ordered response model thus plays a significant role in shaping the implied effect of the independent variable on the probability of the categorical outcomes. In Table 3.2, we normalize the variance in each error distribution to one, to account for the effects of different error distributions on the interpretation of $\widehat{\beta}$. However, this still ignores both some of the power that more flexibledistributions have on estimating the marginal effect of x on each outcome $\{y=j\}$, and the potential impact of changing the values of $\widehat{\alpha}$. Here we present another useful measure which intuitively captures the average effect of changing x on the dependent variables. In order to keep this measure parsimonious, we do not focus on the effect of changing x on the probability of each outcome $\{y=j\}$, but instead we focus on the effect of changing x on $E[y|x]$. We calculate this derivative by summing the marginal changes in probabilistic outcomes as follows:

$$\frac{\delta \widehat{E(y|x)}}{\delta x} = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{J} j \frac{\delta P(\widehat{y_i = j}|x_i)}{\delta x_i} \tag{3.12}$$

In Table 3.4, we will compare estimates of the average value $\frac{\delta E(y|x)}{\delta x}$ which we will refer to as $\hat{\gamma}$. We calculate by estimating the β, α and θ coefficients in Equation 3.3, using coefficients to calculate the derivatives in Equation 3.11 at each point in the sample, and averaging these derivatives over all outcomes and individuals to find the expected value as in Equation 3.12. Notice that $\hat{\gamma}$ is an estimate for the average marginal effect of changing x on total satisfaction in the economy. For example, if the entry of the implied $\hat{\gamma}$ would be an estimate of the expected change in the average level of satisfaction if everyone in the economy were to have their income increase by a dollar. In Table 3.4, we compare the estimates of $\hat{\gamma}$ implied by the MLE ordered response models in Table 3.2.

It is interesting to note that a specific regression having a higher estimate $\widehat{\beta}$ coefficient for independent variable does not necessarily mean it will estimate a higher average marginal effect. For example, notice that in Table 3.2, the SGT has the lowest estimated effect of the religious dummy variable on satisfaction. However, in Table 3.4, the SGT estimates that the religious dummy variable has a greater average marginal effect on satisfaction than do the Laplace, Skewed Laplace or SGED estimates.

Notice that we could construct an average marginal effect measure similar to γ, but instead of measuring the average marginal effect of changing the independent variables on the total satisfaction score reported, we could measure the average marginal effect of changing the independent variables on which satisfaction score is reported. For a given estimator, we would define J average marginal effect measures. These measures would each be a 1 X k vector ϕ_j, for $1 \leq j \leq J$, which would be an estimate of the average value of $\frac{\delta P(y_i=j|x_i)}{\delta x_i}$ across all x's in the sample. We include the tables of the $\widehat{\phi}_j$ vectors for the SGT and Normal estimates in Table 3.5. In Table 3.5 we also include the ratios of the SGT estimates of $\widehat{\phi}_j$ j to the equivalent estimates in the probit estimation in order to help the reader visualize the differences between the two estimators.

The first section of Table 3.5 contains the SGT estimates of ϕ, the second section contains the Normal estimates of ϕ, and the last section contains the ratio of the SGT estimates of ϕ to the Probit estimates of ϕ.

First, it is interesting to note that though the SGT estimates religious status, age, age2, and marital status to all have a smaller average marginal effect on satisfaction (from Table 3.4) than does the Normal, the SGT estimates have a larger marginal effect on whether an individual is in the first two brackets. It appears that the flexibility in the tails of the SGT allows for each of these factors to have a larger marginal effect on whether an individual is in the lowest two satisfaction groups, but a much smaller effect on most of the other brackets. The ratios are negative in the $j = 6$ bracket because the SGT estimates average marginal effects switch signs between the $j = 5$ and the $j = 6$ bracket while the normal estimates switch signs between the $j = 6$ and the $j = 7$ estimates.

Analyzing the ratios of the marginal effect estimates shows some of the power in using flexible error distributions. Using flexible error distributions in an ordered response model not only causes the estimates of β and γ to change significantly from more restrictive distributional assumptions, but it also allows for more flexibility in estimating the average marginal effects on the probability of each outcome occurring.

3.5 Monte Carlo simulations

Overview

To compare the properties of the alternative estimators of b (the coefficient of x) previously discussed we performed Monte Carlo simulations of ordered response models using probit, Laplace, SNormal, GED, SLaplace, SGED, SGT, and logit data generating processes for the error. The generating model for the Monte Carlo simulations is as follows:

$$y_i^* = x_i\beta + \epsilon \; 1 \leq i \leq N \tag{3.13}$$

$$y_i = \begin{cases} 1 & \text{if } y_i^* < \alpha_1 \\ 2 & \text{if } \alpha_1 \leq y_i^* < \alpha_2 \\ 3 & \text{if } \alpha_2 \leq y_i^* < \alpha_3 \\ \vdots & \\ J & \text{if } \alpha_{J-1} \leq y_i^* \end{cases} \tag{3.14}$$

For each of the data generating error distributions, a Monte Carlo simulation is run in which there are $J = 2$possible outcomes (a binary response model), and a simulation in which there are $J = 5$ outcomes. In all of the simulations, β is a $1x1$ vector with a value of [1.5]. The independent variable vector x is thus a $Nx1$ vector that is randomly generated using a normal distribution with a mean of 0 and variance of 1.

Thus we can express Equation 3.13 as

$$y_i^* = 1.5x_i + \epsilon_i \; 1 \leq i \leq N \tag{3.15}$$

In the simulations with 2 possible outcomes, the generating a vector is [1.0]. In the simulations with 5 possible outcomes, the generating a vector is $[-1.0, -0.6, 0.1, 1.0]$. As in Equation 3.2, the a vector contains the cutoff points between the discrete outcomes. As in our empirical results in Section 3.4, we normalize the estimated error distribution to have a mean of 0 and variance of 1. For each distribution and response variable combination, we ran simulations with sample sizes of 1,000, 5,000, 10,000 and 20,000 observations. We completed 10,000 replications of the estimation for each sample size. In each simulation, we estimated an ordered response model assuming each of the error distributions previously listed and investigate the bias and root mean squared error (RMSE) of the corresponding estimated coefficients of the independent variable (β). These comparisons also allow for a comparison of the impact of the number of outcomes, distributional misspecification, and sample size summarized in the next three subsections.

Number of ordered outcomes

The results of two different Monte Carlo simulations are reported in Table 3.6. In comparing the bias and RMSE for the different estimators, increasing the number of ordered outcomes is seen to have a significant impact on how efficiently different error distributions estimate the desired effects. Since the assumed error distribution is a normal, the probit model (Normal) should the most efficient and it is. Increasing the number of outcomes reduces the bias and RMSE for all specifications which include the normal as a special case. Furthermore, the flexible distribution parameter estimates are much closer to the distributional parameters of the normal distribution when there are more outcomes. With 5 outcomes, the SNormal, GED, SGED, and SGT perform similarly to the true distribution. Thus, the flexible distributions are able to approximate the normal, a special case. The Laplace doesn't include the normal as a special case, has the largest bias and RMSE of all of the estimators considered, and actually performs worse as the number of outcomes increases, illustrating the dangers of distributional misspecification.

Impact of distributional misspecification

When the error term distribution is known, deciding which error distribution to use in an ordered response model is a trivial problem. However, the actual error distribution is unlikely to have a known functional form. To show the benefit of using flexible distributions, we consider a case in which the actual error distribution is not a special case of the assumed error distribution. In the Table 3.7, we compare the results of assuming a different error distribution when a logistic error distribution is the generating distribution. We chose this distribution because it is commonly used, is computationally tractable, and it is not a member of the SGT tree.

The Skewed Normal and Skewed Laplace distributions are not reported because they yield almost identical results to the Normal and Laplace distributions respectively. Given that the GED has nearly identical results to the SGED, it appears that having flexibility with λ provides little advantage in matching a symmetric logistic distribution. The SGT has less bias than all of its special cases, and its RMSE is only slightly higher than that of the SGED and GED. In fact, the close agreement between the results for the SGT and the Logistic estimators suggests that the SGT has the flexibility to approximate the Logistic in contrast to the use of the probit or Laplace. Recall the poor performance of the Laplace with a normal error distribution. These two examples illustrate the impact of distributional misspecification and potential benefits of selecting an estimation procedure based on using a flexible probability density function.

Impact of sample size

Though assuming an incorrect error distribution may cause an estimator to be in- consistent, it may still be less biased than a correct, more flexible error distribution with small samples. The simulation results reported in Table 3.8 correspond to an SGT data generating process with significant skewness ($\lambda = 0.4$) and sample sizes of 1,000, 5,000, 10,000, and 20,000 observations.

In the simulations with 1,000 observations, when considering the distribution of $\widehat{\beta}$, the SGT distribution yields a higher RMSE than any of the other distributions. Since the SGT is the data generating error distribution, the SGT estimates are consistent, but their small sample properties are dominated by some of its special cases.

Even for small sample sizes (1,000) the sample bias of the SGT is smaller than the others, but the corresponding tails of the density are thicker than the alternatives considered. As the sample size increases from 1,000, to 5,000, to 10,000, to 20,000, a number of observations can be made [2]. First note that the distribution of the ordered probit and ordered logit estimators are relatively invariant to the sample size with a relatively large bias, suggesting that they are inconsistent. The bias and RMSE for the GED and SGED decrease very slowly. The bias and the RMSE of the SGT specification decrease quite rapidly, with the RMSE roughly decreasing as $\frac{1}{\sqrt{n}}$ where n is the sample size. The excellent performance of the Laplace and SLaplace estimators is somewhat surprising, especially considering their poor performance with a Logistic distribution. Notice that the Laplace estimator yields lower RMSE than the SGT for sample sizes of 5,000 or smaller, and the SLaplace yields lower RMSE than the SGT for sample sizes 10,000 or smaller. However, the RMSE and bias decrease at a much slower rate for the Laplace and SLaplace than they do for the SGT, so the Skewed Generalized T has lower bias and RMSE than its special cases in the 20,000 observation case.

3.6 Conclusion

In this chapter, we introduce a generalized ordered probit model which allows the error distribution to take different forms. We present an empirical application of the generalized model and analyze the impact of different error distributions on the estimated results. We found that the Skewed Generalized T gave results that differed significantly from the results of assuming a Normal, Laplace, Logistic, Skewed Normal, GED, Skewed Laplace, or SGED error distribution. We then analyzed the results of Monte Carlo simulations, which indicated that using flexible distributions is more advantageous when there are more ordered outcomes. The Monte Carlo simulations also indicated that more flexible error distributions can be more efficient than less flexible ones when the distributional assumptions are incorrect. The SGT seemed to have the flexibility to accurately model distributions which were not special cases, such as the logit.

[2] The Monte Carlo approximations with four cutoffs and 10,000 observations collectively took 3,677 processor-hours. With 64 of these approximations, each having 10,000 replications, the implied average computational time for an ordered response model with 10,000 observations and four cutoffs is 20.7 seconds.

3.7 Tables and Figures

Table 3.1: World Values Survey Data Summary

Variable	Mean	Std. Dev.	Min	Max
satisfaction	5.34	2.60	0	9
religious	0.83	0.36	0	1
male	0.49	0.50	0	1
income	4.53	2.45	1	10
age	39.68	15.33	15	99
married	0.65	0.48	0	1
unemployed	0.12	0.33	0	1
high school	0.57	0.49	0	1
Observations	104,437			

Notes: Summary statistics for waves 1-5 of the World Values Survey (The waves begin in 1981 and end in 2008).

Table 3.2: World Values Survey Empirical Results

	Normal	Logit	Laplace	SLaplace	SGED	SGT
Religious	0.1286***	0.1279***	0.0988***	0.0869***	0.0961***	0.0873***
Std. Errs	*(0.0091)*	*(0.0085)*	*(0.0064)*	*(0.0060)*	*(0.0068)*	*(0.0062)*
male	−0.0347***	−0.0277***	−0.0217***	−0.0152***	−0.0181***	−0.0131***
	(0.0064)	*(0.0061)*	*(0.0047)*	*(0.0044)*	*(0.0049)*	*(0.0045)*
income	0.0798***	0.0788***	0.0656***	0.0690***	0.0752***	0.0683***
	(0.0014)	*(0.0013)*	*(0.0011)*	*(0.0010)*	*(0.0011)*	*(0.0010)*
age	−0.0266***	−0.0257***	−0.0206***	−0.0180***	−0.0201***	−0.0187***
	(0.0012)	*(0.0011)*	*(0.0009)*	*(0.0008)*	*(0.0009)*	*(0.0008)*
age2	2.82E-04***	2.74E-04***	2.20E-04***	1.84E-04***	2.06E-04***	1.91E-04***
	(1.27E-05)	*(1.22E-05)*	*(9.71E-06)*	*(9.08E-06)*	*(1.00E-05)*	*(9.21E-06)*
married	0.1001***	0.0923***	0.0597***	0.0631***	0.0723***	0.0674***
	(0.0074)	*(0.0070)*	*(0.0055)*	*(0.0051)*	*(0.0057)*	*(0.0052)*
unemployed	−0.0003	−0.0463***	−0.0269***	−0.0129*	−0.0103	−0.0152**
	(0.0099)	*(0.0093)*	*(0.0071)*	*(0.0067)*	*(0.0075)*	*(0.0069)*
high school	0.0019	0.0103	0.0225***	0.0233***	0.0241***	0.0228***
	(0.0067)	*(0.0064)*	*(0.0049)*	*(0.0047)*	*(0.0052)*	*(0.0048)*
Loglikelihood	−230011.4	−229823.8	−229710.8	−229351.4	−229320.5	−229311.8
BIC	460115.3	459740.0	459514.1	458806.8	458756.6	458750.7
Dist. Param						
λ	0	NA	0	0.2718	0.2740	0.2815
p	2	NA	1	1	1.1385	1.3636
q	Inf	NA	Inf	Inf	Inf	4.4713

Notes: Results usi4g increasingly more flexible error distributional assumptions. ***, **, and * respectively mean significant at the .01, .05, and .1 level. In the distributional parameters, 'inf' denotes a limiting case as the specified parameter approaches inf. The SLaplace, SGED, and SGT specifications involve 1, 2, and 3 estimated distributional parameters, respectively.

Table 3.3: LR tests

	Normal vs SGT	Laplace vs SGT	SLaplace vs SGT	SGED vs SGT
LR Test	1399.2	798.1	79.3	17.5
P-value	<1E-10	<1E-10	<1E-10	0.00003

Notes: LR tests between various distributions used in table 3.2

Table 3.4: Empirical estimates of $\hat{\gamma}$

	Normal	Laplace	SLaplace	SGED	SGT
religious	0.34979	0.24875	0.23344	0.26603	0.28241
male	−0.09447	−0.05470	−0.04085	−0.05006	−0.04226
income	0.21710	0.16509	0.18539	0.20814	0.22092
age	−0.07232	−0.05175	−0.04833	−0.05570	−0.06041
age2	0.00077	0.00055	0.00050	0.00057	0.00062
married	0.27238	0.15034	0.16945	0.20018	0.21799
unemployed	−0.00070	−0.06767	−0.03480	−0.02857	−0.04906
high school	0.00530	0.05665	0.06258	0.06675	0.07380

Notes: Estimates of the average value of $\frac{\delta E(y|x)}{\delta x}$.

Table 3.5: Empirical estimates of $\hat{\phi}_j$

	religious	male	income	age	age2	married	Unemployed	high school
SGT								
j = 1	−0.01088	0.00163	−0.00851	0.00233	−0.00002	−0.0084	0.00189	−0.00284
j = 2	−0.00845	0.00126	−0.00661	0.00181	−0.00002	−0.00652	0.00147	−0.00221
j = 3	−0.01157	0.00173	−0.00905	0.00248	−0.00003	−0.00893	0.00201	−0.00302
j = 4	−0.01016	0.00152	−0.00794	0.00217	−0.00002	−0.00784	0.00176	−0.00265
j = 5	−0.01007	0.00151	−0.00788	0.00215	−0.00002	−0.00777	0.00175	−0.00263
j = 6	0.00206	−0.00031	0.00161	−0.00044	−0.00000	0.00159	−0.00036	0.00054
j = 7	0.00869	−0.0013	0.0068	−0.00186	0.00002	0.00671	−0.00151	0.00227
j = 8	0.01501	−0.00225	0.01174	−0.00321	0.00003	0.01158	−0.00261	0.00392
j = 9	0.0111	−0.00166	0.00869	−0.00238	0.00002	0.00857	−0.00193	0.0029
j = 10	0.01427	−0.00213	0.01116	−0.00305	0.00003	0.01101	−0.00248	0.00373
Normal								
	religious	male	income	age	age2	married	Unemployed	high school
j = 1	−0.00623	0.00168	−0.00387	0.00129	−0.00001	−0.00485	0.00001	−0.00009
j = 2	−0.00802	0.00217	−0.00498	0.00166	−0.00002	−0.00625	0.00002	−0.00012
j = 3	−0.01319	0.00356	−0.00819	0.00273	−0.00003	−0.01027	0.00003	−0.0002
j = 4	−0.01361	0.00368	−0.00845	0.00281	−0.00003	−0.0106	0.00003	−0.00021
j = 5	−0.02319	0.00626	−0.01439	0.00479	−0.00005	−0.01806	0.00005	−0.00035
j = 6	−0.00523	0.00141	−0.00324	0.00108	−0.00001	−0.00407	0.00001	−0.00008
j = 7	0.00538	−0.00145	0.00334	−0.00111	0.00001	0.00419	−0.00001	0.00008
j = 8	0.02232	−0.00603	0.01385	−0.00461	0.00005	0.01738	−0.00004	0.00034
j = 9	0.0205	−0.00554	0.01273	−0.00424	0.00004	0.01597	−0.00004	0.00031
j = 10	0.02126	−0.00574	0.0132	−0.0044	0.00005	0.01655	−0.00004	0.00032
SGT/Normal								
	religious	male	income	age	age2	married	unemployed	high school
j = 1	1.74747	0.96808	2.20246	1.80797	1.74231	1.73217	152.09331	30.11847
j = 2	1.05364	0.58371	1.32797	1.09012	1.05053	1.04441	91.70479	18.15996
j = 3	0.8773	0.48601	1.10572	0.90767	0.8747	0.86962	76.35651	15.12059
j = 4	0.74615	0.41336	0.94042	0.77198	0.74395	0.73962	64.94207	12.86023
j = 5	0.43429	0.2406	0.54737	0.44933	0.43301	0.43049	37.79927	7.48525
j = 6	−0.39452	−0.21856	−0.49724	−0.40818	−0.39335	−0.39107	−34.33749	−6.79972
j = 7	1.61476	0.89456	2.03519	1.67066	1.60999	1.60062	140.54232	27.83107
j = 8	0.67237	0.37249	0.84743	0.69565	0.67038	0.66648	58.5203	11.58855
j = 9	0.54156	0.30002	0.68257	0.56031	0.53996	0.53682	47.13543	9.33405
j = 10	0.67106	0.37176	0.84578	0.69429	0.66908	0.66518	58.40649	11.56602

Notes: Estimates of the average marginal effect of changing the independent variables on which satisfaction score is reported in Table 3.2

Table 3.6: Analyzing the effect of more potential outcomes

	Normal	Laplace	SNormal	GED	SGED	SGT
2 Outcomes						
Bias	0.00176	−0.15152	0.00156	0.00106	0.00104	−0.01038
RMSE	0.04508	0.15819	0.06478	0.04558	0.08311	0.07660
p	2	1	2	2.04172	2.09624	2.24327
Std. Dev				(0.28196)	(0.47834)	(1.52109)
λ	0	0	0.00390	0	0.03617	0.00655
			(0.10133)		(0.17954)	(0.14438)
q	∞	∞	∞	∞	∞	1.571E + 08
						(4.566E + 08)
5 Outcomes						
	Normal	Laplace	SNormal	GED	SGED	SGT
Bias	0.00096	−0.27271	0.00081	0.00058	0.00035	−0.00051
RMSE	0.02929	0.27417	0.02930	0.03095	0.03098	0.03114
p	2	1	2	2.00442	2.00325	2.02594
				(0.12415)	(0.12425)	(0.14717)
λ	0	0	0.00055	0	0.00049	−0.00536
			(0.03644)		(0.03671)	(0.02963)
q	Inf	Inf	Inf	Inf	Inf	3.990E + 07
						(9.564E + 07)

Notes: Results from Monte Carlo simulations. In simulations with two possible outcomes, the generating α vector is [1.0]. In the simulations with 5 possible outcomes, the generating α vector is [-1.0, -0.6, 0.1, 1.0]. We normalize the estimated error distribution to have a mean of 0 and a variance of 1. We performed 10,000 simulations using a sample size of 5,000 and a $\beta = 1.5$. The means of the distributional parameter estimates are included, along with the parameter estimate standard deviations in parenthesis. Inf denotes the limiting case where the specified parameter approaches 1.

Table 3.7: Impact of distributional misspecification

	Normal	Laplace	GED	SGED	SGT	Logistic
Bias	0.01964	−0.09679	0.00773	0.00764	0.00175	0.00172
RMSE	0.03742	0.10156	0.03397	0.03396	0.03547	0.03117

Logit data generating process with 5 outcomess, 5000 observations. $\beta = 1.5$

Table 3.8: Impact of sample size

	Normal	Laplace	Logit	GED	Slaplace	SGED	SGT
1,000 observations							
Bias	0.21168	0.09608	0.22427	0.12983	0.06647	0.11176	−0.04605
RMSE	0.22856	0.12537	0.23879	0.16565	0.10487	0.15253	0.26363
5,000 observations							
	Normal	Laplace	Logit	GED	Slaplace	SGED	SGT
Bias	0.20254	0.08937	0.21799	0.12257	0.06052	0.10532	−0.01245
RMSE	0.20621	0.09623	0.22101	0.13091	0.07041	0.11498	0.10339
10,000 observations							
	Normal	Laplace	Logit	GED	Slaplace	SGED	SGT
Bias	0.20179	0.08866	0.21736	0.12207	0.05993	0.10481	−0.00627
RMSE	0.20364	0.09212	0.21885	0.12627	0.06502	0.10972	0.06988
20,000 observations							
	Normal	Laplace	Logit	GED	Slaplace	SGED	SGT
Bias	0.20154	0.08835	0.21710	0.12229	0.05966	0.10474	−0.00250
RMSE	0.20251	0.09010	0.21784	0.12438	0.06226	0.10722	0.04775

Notes: Each regression has 5 outcomes and an SGT generating distribution with parameters $\gamma = .4$, $p = 1.7$, and $q = 2.6$. The true value of β is 1.5.

Figure 3.1: The SGT distribution tree

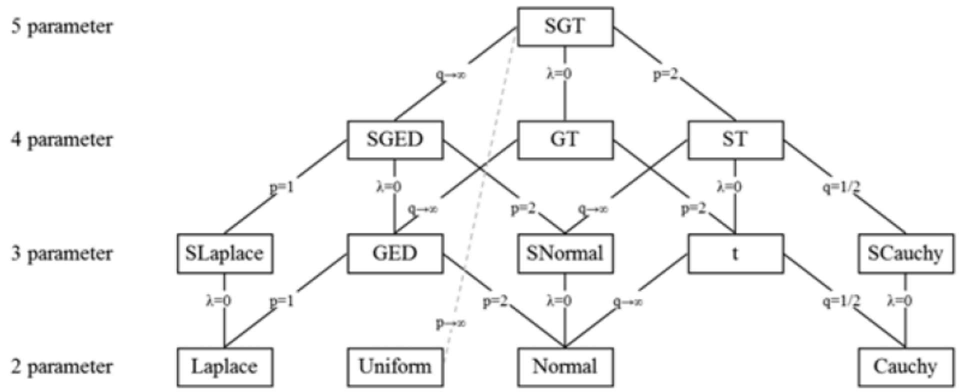

Adapted from Hansen et al. (2010)

The relationship between the SGT and many of its special and limiting cases. The parameter m is a location parameter, λ controls skewness, and p and q are positive shape parameters that determine peakedness and kurtosis.

Conclusion

This dissertation explores two government subsidies and presents a model to improve causal effect estimation for categorical, or binned variables. I have found that the nonprofit tax break does not significantly affect workers' wages in the largest nonprofit industries, health and education. I have also shown that in many industries, nonprofit and for-profit firms are similar in many ways except for simply the tax break. This warrants further research into understanding whether the tax break is correcting a free market aberration in some industries such as healthcare, given that for-profit firms make enough revenues to exist profitably.

In my second chapter I have found encouraging evidence regarding the efficacy of the Earned Income Tax Credit. Families do appear to spend the tax credit upon receipt, indicating the credit is helping families fulfill a consumption need. Additionally, at least in the used car market, dealers are not altering their prices during the issuance month, despite increased demand. These two pieces of evidence suggest that the structure of the tax credit and the timing of its receipt result in intended policymaker outcomes.

My third chapter supplies applied researchers with a tractable model that allows them to more consistently estimate causal effects using binned, publicly-available, government data. This model is especially useful for the type of aggregated data issued by the IRS.

Bibliography

Abowd, John M, Francis Kramarz, and David N Margolis. 1999. “High wage workers and high wage firms.” *E*conometrica, 67(2): 251–333.

Arrow, Kenneth J. 1978. “Uncertainty and the welfare economics of medical care.” In *U*ncertainty in economics.: Elsevier, 345–375.

Barrow, Lisa, and Leslie McGranahan. 2000. “The effects of the earned income credit on the seasonality of household expenditures.” *N*ational Tax Journal 1211–1243.

Baumol, William J, and William G Bowen. 1965. “On the performing arts: the anatomy of their economic problems.” *T*he American economic review, 55(1/2): 495–502.

Bishow, John L, and Kristen A Monaco. 2016. “Nonprofit pay and benefits: estimates from the National Compensation Survey.” *M*onthly Lab. Rev., 139, p. 1.

Borjas, George J, HE Frech III, and Paul B Ginsburg. 1983. “Property rights and wages: The case of nursing homes.” *J*ournal of Human Resources 231–246.

Bound, John, and Alan B Krueger. 1991. “The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?” *J*ournal of Labor Economics, 9(1): 1–24.

Card, David. 1996. “The effect of unions on the structure of wages: A longitudinal analysis.” *E*conometrica: Journal of the Econometric Society 957–979.

Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women.” *T*he Quarterly Journal of Economics, 131(2): 633–686.

Chen, Songnian, and Shakeeb Khan. 2003. “Rates of convergence for estimating regression coefficients in heteroskedastic discrete response models.” *J*ournal of Econometrics, 117(2): 245–278.

Chiswick, Barry R. 1991. “Speaking, reading, and earnings among low-skilled immigrants.” *J*ournal of labor economics, 9(2): 149–170.

Coppejans, Mark. 2007. “On efficient estimation of the ordered response model.” *Journal of Econometrics*, 137(2): 577–614.

Donat, Francesco, and Giampiero Marra. 2017. “Semi-parametric bivariate polychotomous ordinal regression.” *Statistics and Computing*, 27(1): 283–299.

Donat, Francesco, and Giampiero Marra. 2018. “Simultaneous equation penalized likelihood estimation of vehicle accident injury severity.” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 67(4): 979–1001.

Feldstein, Martin S. 1971. “Hospital cost inflation: A study of nonprofit price dynamics.” *The American Economic Review*, 61(5): 853–872.

Frank, Robert H. 1996. “What price the moral high ground?” *Southern Economic Journal* 1–17.

Glaeser, Edward L. 2002. “The governance of not-for-profit firms.”Technical report, National Bureau of Economic Research.

Goddeeris, John H. 1988. “Compensating differentials and self-selection: An application to lawyers.” *Journal of Political Economy*, 96(2): 411–428.

Goodman-Bacon, Andrew, and Leslie McGranahan. 2008. “How do EITC recipients spend their refunds?” *Economic Perspectives*, 32(2): .

Greene, William H, and David A Hensher. 2010. *Modeling ordered choices: A primer.*: Cambridge University Press.

Hansen, Bruce E. 1994. “Autoregressive conditional density estimation.” *International Economic Review* 705–730.

Hansmann, Henry B. 1980. “The role of nonprofit enterprise.” *The Yale law journal*, 89(5): 835–901.

Hirsch, Barry T, David A Macpherson, and Anne Preston. 2017. “Nonprofit Wages: Theory and Evidence.”

Hodge, Andrew, and Sriram Shankar. 2014. “Partial effects in ordered response models with factor variables.” *Econometric Reviews*, 33(8): 854–868.

Holtmann, Alphonse G, and Todd L Idson. 1993. “Wage determination of registered nurses in proprietary and nonprofit nursing homes.” *Journal of Human Resources* 55–79.

Horwitz, Jill R, and Austin Nichols. 2007. “What do nonprofits maximize? Nonprofit hospital service provision and market ownership mix.”Technical report, National Bureau of Economic Research.

Jones, Daniel B. 2015. "The supply and demand of motivated labor: When should we expect to see nonprofit wage gaps?" *L*abour Economics, 32 1–14.

Kerman, Sean C, and James B McDonald. 2013. "Skewness–kurtosis bounds for the skewed generalized T and related distributions." *S*tatistics & Probability Letters, 83(9): 2129–2134.

Klein, Roger W, and Robert P Sherman. 2002. "Shift restrictions and semiparametric estimation in ordered response models." *E*conometrica, 70(2): 663–691.

LaLumia, Sara. 2013. "The EITC, tax refunds, and unemployment spells." *A*merican Economic Journal: Economic Policy, 5(2): 188–221.

Leete, Laura. 2001. "Whither the nonprofit wage differential? Estimates from the 1990 census." *J*ournal of labor Economics, 19(1): 136–170.

Leete, Laura. 2006. "Work in the Nonprofit Sector." In *T*he Nonprofit Sector. eds. by Walter W. Powell, and Richard Steinberg, Oxford: Yale University Press, , Chap. 7 159–179.

Lemp, Jason D, Kara M Kockelman, and Avinash Unnikrishnan. 2011. "Analysis of large truck crash severity using heteroskedastic ordered probit models." *A*ccident Analysis & Prevention, 43(1): 370–380.

Lewbel, Arthur et al. 2003. "Ordered response threshold estimation." *U*npublished working paper, Boston College.

Lindblad, Marilyn., "Do Nonprofit Organizations Pay Property Taxes?" *S*mall Business-Chron. com.

Matzkin, Rosa L. 1993. "Nonparametric identification and estimation of polychotomous choice models." *J*ournal of Econometrics, 58(1-2): 137–168.

McDonald, James B, and Whitney K Newey. 1988. "Partially adaptive estimation of regression models via the generalized t distribution." *E*conometric theory, 4(3): 428–457.

Mocan, H Naci, and Erdal Tekin. 2003. "Nonprofit sector and part-time work: An analysis of employer-employee matched data on child care workers." *R*eview of Economics and Statistics, 85(1): 38–50.

Mora, Juan, and Ana I Moro-Egido. 2008. "On specification testing of ordered discrete choice models." *J*ournal of Econometrics, 143(1): 191–205.

Newhouse, Joseph P. 1970. "Toward a theory of nonprofit institutions: An economic model of a hospital." *T*he American economic review, 60(1): 64–74.

Nichols, Austin, and Jesse Rothstein. 2015. "The earned income tax credit." In *Economics of Means-Tested Transfer Programs in the United States, Volume 1*.: University of Chicago Press, 137–218.

Pauly, Mark, and Michael Redisch. 1973. "The not-for-profit hospital as a physicians' cooperative." *The American Economic Review*, 63(1): 87–99.

Preston, Anne E. 1989. "The nonprofit worker in a for-profit world." *Journal of labor economics* 438–463.

Pudney, Stephen, and Michael Shields. 2000. "Gender, race, pay and promotion in the British nursing profession: estimation of a generalized ordered probit model." *Journal of Applied Econometrics*, 15(4): 367–399.

Rigby, Robert A, and D Mikis Stasinopoulos. 2005. "Generalized additive models for location, scale and shape." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 54(3): 507–554.

Rose-Ackerman, Susan. 1996. "Altruism, nonprofits, and economic theory." *Journal of economic literature*, 34(2): 701–728.

Ruhm, Christopher J, and Carey Borkoski. 2003. "Compensation in the nonprofit sector." *Journal of Human Resources*, 38(4): 992–1021.

Salamon, Lester M, and S Wojciech Sokolowski. 2005. "Nonprofit organizations: New insights from QCEW data." *Monthly Lab. Rev.*, 128, p. 19.

Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up inequality." *The Quarterly journal of economics*, 134(1): 1–50.

Theodossiou, Panayiotis. 1998. "Financial data and the skewed generalized t distribution." *Management Science*, 44(12-part-1): 1650–1661.

Ventry Jr, Dennis J. 2000. "The collision of tax and welfare politics: The political history of the Earned Income Tax Credit, 1969–99." *National Tax Journal* 983–1026.

Weisbrod, Burton Allen. 1975. "Toward a theory of the voluntary non-profit sector in a three-sector economy."

Wood, Simon N, Natalya Pya, and Benjamin Säfken. 2016. "Smoothing parameter and model selection for general smooth models." *Journal of the American Statistical Association*, 111(516): 1548–1563.

Yee, Thomas W, and CJ Wild. 1996. "Vector generalized additive models." *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(3): 481–493.

Zweimüller, Josef, and Rudolf Winter-Ebmer. 1994. "Gender wage differentials in private and public sector jobs." *Journal of Population Economics*, 7(3): 271–285.