Essays on Public Economics

by

Tzu-Ting Yang

B.A., National Taiwan University, 2005
M.A., National Taiwan University, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Economics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

April 2015

Abstract

This dissertation applies various program evaluation techniques to examine both the intended and unintended consequences of government spending and regulation that affect family labor supply, children's healthcare utilization, and household saving behavior. Chapter 2 of this dissertation exploits the large and anticipated cash influx in the first quarter of the calendar year induced by the Earned Income Tax Credit (EITC) to estimate the causal effect of the receipt of a cash transfer on the timing of family labor supply. I find that income seasonality caused by EITC receipt induces changes in the intra-year labor supply patterns of married women. In contrast, the receipt of the EITC does not affect the timing of the labor supply of married men and single women. The subgroup analysis implies that my results are mainly driven by those from liquidity-constrained families and those who are secondary earners within their families. Chapter 3 exploits a sharp increase in patient cost-sharing at age 3 in Taiwan that results from young children "aging out" of the cost-sharing subsidy to examine the causal effect of cost sharing on the demand for young children's healthcare. It shows that the increase in the level of patient cost sharing at the 3rd birthday significantly reduces utilization of outpatient care. However, the utilization of inpatient care for young children *does not* respond to a change in cost sharing at the 3rd birthday. Chapter 4 exploits workplace pension reform in Taiwan to estimate the casual effect of workplace pension provision on the household saving rate. It shows that pension reform significantly reduces the prime-age (20–50) household saving rate by between 2.06 percentage points and imply that the degree of substitutability between workplace pensions and saving is about -0.50 to -0.60.

Preface

Chapter 3 uses data from Taiwan's National Health Insurance Research Database (NHIRD). Ethics approval under the project title *Patient Cost-Sharing and Health Care Utilization in Early Childhood: Evidence from a Regression Discontinuity Design* was obtained through the Behavioural Research Ethics Board of the University of British Columbia (H14-00869).

Chapter 3 *Patient Cost-Sharing and Health Care Utilization in Early Childhood: Evidence from a Regression Discontinuity Design* of this dissertation is a joint work with Hsing-Wen Han and Hsien-Ming Lien. Dr. Hsing-Wen Han is an Associate Professor from the Department of Accounting at the Tamkang University. Dr. Hsien-Ming Lien is a Professsor from the Department of Public Finance at the National Chengchi University. I was highly involved throughout every stage of the research: collecting and preparing data, designing empirical models, carrying out estimation, organizing and presenting results, writing and editing the manuscript.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I am extremely grateful to my supervisors Joshua Gottlieb and Kevin Milligan for their countless hours of advice and supervision. Their sharp sense in economics and strict academic standard greatly improved my work. In addition, both of them were generous in their time and gave me timely help throughout the job market "year". Without them, I would not have been able to get my dream job. I would also like to thank my dissertation committee members Thomas Lemieux and Marit Rehavi for their insightful suggestions and much encouragement. I thank my university (external) examiners: Mukesh Eswaran, Thomas Davidoff, and Ted McDonald for carefully reading my thesis and making useful comment. I also benefited greatly from talking to many VSE faculty members Patrick Francois, Matilde Bombardini, Nicole Fortin, David Green, Florian Hoffmann, Amartya Lahiri, Craig Riddell, and Giovanni Gallipoli. I also thank my coauthors: Hsien-Ming Lien and Hsing-Wen Han for their help and support. I am also indebted to Maureen Chin for her invaluable help in dealing with the administrative work throughout the program, especially during the job market.

I want to thank my outstanding colleagues and friends at UBC, Alexandre Corhay, Francis Michaud, Zhengfei Yu, Zhe Chen, Yawen Liang, Bipasha Maity, Guidon Fenig, Hugo Jales, Nathan Canen, and Sharon Hu for their support. I thank my colleagues in the public finance reading group, Alix Duhaime-Ross, Oscar Becerra, Timea Laura Molnar, Daniel Shack, and Lori Timmins, for their generous feedback and insightful discussion in the last two years of the program.

I benefited greatly from sharing the daily life with residents of St John's College (SJC) at dining time. Many of them have become my lifetime friends. I also want to thank my friends at Taiwanese Graduate Student

Association in Vancouver (TGSA). Joining TGSA activities greatly cures my homesick.

I also wish to acknowledge classmates and faculty members at the National Taiwan University, particularly Ming-Ching Luoh, who encouraged me to study abroad.

Finally, I would like to express my deepest gratitude to my parents, to my brother and to Yi-Ling Lin for their unconditional support, encouragement and love.

Dedication

獻給親愛的父母, 弟弟與妻子
For my parents, brother, and wife

Chapter 1

Introduction

This dissertation applies various program evaluation techniques to examine both the intended and unintended consequences of government spending and regulation that affect family labor supply, children's healthcare utilization, and household saving behavior.

Dissertation Outline

The second chapter, *"Family Labor Supply and the Timing of Cash Transfers: Evidence from the Earned Income Tax Credit"* provides new evidence on how families adjust their labor supply in response to the receipt of an anticipated cash transfer. In particular, I exploit the unique disbursement timing and benefit rules of the Earned Income Tax Credit (EITC) to assess the effect of the receipt of a cash transfer on the timing of family labor supply. My results show that income seasonality caused by EITC receipt leads to changes in the intra-year labor supply patterns of married women. On average, receiving a $1,000 EITC payment significantly reduces the proportion of married women who work, by 1.6 percentage points, in the month in which the EITC is received. The income elasticity of labor supply for married women based on this estimate is around -0.06. In contrast, the receipt of the EITC does not affect the timing of the labor supply of married men and single women. The subgroup analysis suggests families might reduce the labor supply of secondary earners in response to receiving an anticipated EITC payment. In addition, My results suggest that the presence of liquidity constraints and myopia could be important reasons for my findings.

Healthcare for young children is highly subsidized in many public health insurance programs around the world. However, the existing literature lacks evidence on how the demand for young children's healthcare reacts to these medical subsidy policies. The third chapter *"Patient Cost-Sharing and Healthcare Utilization in Early Childhood: Evidence from a Regression Discontinuity Design"* (joint with Hsing-Wen Han and Hsien-Ming Lien) exploits a sharp increase in patient cost sharing — the share of healthcare costs the patient must pay out of their own pocket — at age 3 in Taiwan, resulting from young children "aging out" of the cost-sharing subsidy. This price shock on the 3rd birthday allows us to use a regression discontinuity design to examine the causal effect of cost sharing on the demand for young children's healthcare by comparing the utilization of healthcare for young children right before and right after their 3rd birthday. Our results show that the increase in the level of patient cost sharing at the 3rd birthday significantly reduces total outpatient expenditure. The implied price elasticity of outpatient expenditure is around -0.10. However, the demand for inpatient care for young children *does not* respond to a change in cost sharing at the 3rd birthday even though the price variation is much larger. This result implies that providing full insurance coverage for children's inpatient care can substantially reduce the financial risk for the households but does not induce excessive utilization of inpatient care.

Population aging causes financial imbalance in pay-as-you-go public pension programs. To remedy this problem but also ensure the adequacy of retirement savings for employees, many countries complement or substitute for public pensions by regulating workplace pensions. The fourth chapter *"The Effect of Workplace Pensions on Household Saving: Evidence from a Natural Experiment in Taiwan"* is the first to utilize a national pension policy change as a natural experiment to identify the impact of employer-sponsored pensions on voluntary household saving. Specifically, I evaluate the response in household saving to a workplace pension reform in Taiwan that has mandated, since 2005, that all private-sector employers contribute at least 6% of wages to employees' individual pension accounts monthly. I use the workers in the unaffected sector as a comparison group and employ a difference-in-differences method to estimate the impact of the reform on the household saving rate. My estimates suggest that making private pensions mandatory significantly reduces the prime-age (20–50) household saving rate by between 2.06 and 2.45 percentage points and imply that the degree of substitutability between workplace pensions and saving is about -0.50 to -0.60. Since workplace pensions only partially offset household saving, a mandatory workplace pension policy could effectively raise employees' retirement wealth.

Chapter 2

Family Labor Supply and the Timing of Cash Transfers: Evidence from the Earned Income Tax Credit

2.1 Introduction

Do households adjust their behavior in response to receiving expected income payments? This question is crucial to understanding households' behavior and analyzing aspects of government policies. For example, the answer to this question has important implications for the design of welfare programs, especially for determining the payment frequency of welfare benefits. If the benefit recipient's behavior is sensitive to the receipt of income, then more frequent payments could improve policy by helping recipients to smooth out their consumption. On the other hand, the effectiveness of short-run fiscal policies, such as temporary rebates during a recession, largely depends on how people adjust their behavior after receiving payments. The central implication of the life cycle model with a perfect credit market is that consumption behavior should not respond to predictable changes in income. A growing empirical literature tests this claim by examining whether the timing of the receipt of income is associated with the timing of household spending. Most prior studies find that families increase their spending right after they receive expected income payments, such as a public pension (Stephens, 2003; Stephens and Unayama, 2011), temporary rebate

(Johnson et al., 2006, 2013; Parker, 2014) or income tax refund (Souleles, 1999). These findings generally offer evidence against the theory.[1]

This paper deviates from the previous studies by considering an economically important but seldom addressed question — Do families change their labor supply in response to the receipt of an anticipated cash transfer?[2] Empirically investigating the labor supply response has important implications for both economic theory and public policy. On the theoretical side, such investigation examines one central prediction of the life cycle model of labor supply: Any anticipated income changes should not affect labor supply behavior. When families are informed of a future income change, they can adjust their labor supply (i.e. leisure consumption) in advance through borrowing and saving. Thus, there should be no change in labor supply at the time when the income change is experienced. Several recent studies (e.g. Saez, 2003; Looney and Singhal, 2006) use this prediction to assume away income effects associated with anticipated changes in tax rates when estimating the intertemporal substitution elasticity of labor supply. However, its validity is questionable given the vast evidence on the household spending response to predictable changes in income. Regarding policy, such investigation helps us better understand how the timing of welfare bene-

[1] Johnson et al. (2013) exploit the random timing of the receipt of tax rebates (i.e. US economic stimulus payments) to examine the spending response to the receipt of income. They find that households spend 60% of the tax rebate within three months of receiving it. Stephens and Unayama (2011) use the Japanese public pension which is distributed every three months; they find that household spending closely follows the disbursement patterns of the public pension benefit. Stephens (2003) finds that households spend more on weekly nondurable consumption upon receiving the monthly US Social Security check. Souleles (1999) provides evidence that households in the US spend around 30% to 60% of an income tax refund within one quarter of receiving it. However, not all tests of the life cycle model that examine the spending response to the receipt of income find evidence contradicting consumption-smoothing behavior. Hsieh (2003) utilizes Alaska's annual oil revenue dividend payment, which is equal to two-thirds of the monthly pre-tax household income. He finds that Alaskan households do not change their spending when they receive this payment. Paxson (1993) finds that Thai households can smooth out their spending over a striking degree of variation in seasonal income. Jappelli and Pistaferri (2010) provide an excellent review of this strand of the literature.

[2] Previous studies examining the validity of the life cycle model usually abstract from the labor supply decision and implicitly assume that saving and borrowing are the only ways to smooth out household spending.

fit payments affects household behavior. Past policy debates have mainly focused on the impact of the timing on household spending. My results will show that the timing of cash transfers also matters for the family labor supply decision.

I investigate the above issue by assessing the immediate labor supply response to the receipt of the Earned Income Tax Credit (EITC), a refundable tax credit that subsidizes the earnings of the working poor.[3] The EITC is the largest cash transfer program for low-income families in the US.[4] There are two features of the EITC that make it an interesting case to study the issue I want to address. First, it is widely known and highly anticipated by the recipients. Previous studies suggest that most EITC recipients know what their EITC refund will be before filing their taxes (Chetty and Saez, 2013; Romich and Weisner, 2000).

Second, the EITC could be the single largest cash transfer that many of the working poor will receive during the year. The payment is fairly large relative to the recipients' family income.[5] The average amount of EITC for eligible families is around $2,000 and can account for one month of family income.[6] For some families, it can comprise as much as 45% of their annual income. In addition, most EITC recipients receive their credit in the form of a one-time lump-sum payment within a narrow time frame.[7] Figure 2.3

[3]The EITC is fully refundable. That is, any excess credit beyond a family's income tax liability will be paid in the form of a tax refund. Over 90% of the value of the EITC is delivered in the form of tax refunds, as opposed to serving to reduce tax liabilities (McGranahan and Schanzenbach, 2014).

[4]The structure of the US welfare system has undergone a substantial change in the past three decades. The federal income tax system has become a major policy tool for providing cash assistance to low-income families with children (Eissa and Hoynes, 2011). The expansion of the EITC accounts for most of this dramatic transformation. In 2011, the federal government spent $61 billion on the EITC, substantially more than the $27.1 billion spent on Temporary Assistance for Needy Families (TANF), the flagship cash transfer program in the US, while the two programs were of a similar size in 1994 (see Figure 2.1).

[5]The amount of the EITC largely depends on the family's income, the number of children, and the marital status of the taxpayer in the previous year. I will discuss the EITC benefit rules in detail in Section 2.2.

[6]This number is based on the EITC payment schedule during my sample period (1997–2012).

[7]Recipients formerly had the option of receiving installments of their expected credit

shows that over half of EITCs are paid in the month of February.[8] This concentrated delivery of cash transfers induces a large variation in families' disposable income across 12 months.[9] I use it to examine how family labor supply reacts to the receipt of anticipated cash transfers by providing the first evidence on how the EITC payment timing affects the timing of the family labor supply.[10]

Since most EITC recipients are low-income families with children, the primary concern over relying on variation in payment timing is that my estimates may simply reveal the intra-year labor supply patterns of specific demographic groups, such as those with low incomes or those who have children, rather than reflecting the impact of receiving the EITC payment. I deal with this concern in two ways.

First, I conduct triple differences estimations by using a comparison group of individuals, such as those with children but with incomes just above the EITC range or those without children but incomes within the EITC range, who are similar to those in the treatment group in many ways but receive much smaller EITC payments to control any confounding effects unrelated to the receipt of the EITC. Note that membership in the treat-

on a monthly basis in the calendar year prior to tax filing (Advance EITC). This advance payment option has been unavailable since 2011 due to a very low take-up rate. I will discuss this issue in Section 2.2.

[8]This is because recipients need to file their taxes first and will then obtain their refund a few weeks later. The US government usually opens the tax window in the middle of January and low-income taxpayers tend to file their taxes early, meaning that most tax refunds for low-income families are issued in February.

[9]Previous studies find strong evidence that the timing of the household spending of EITC recipients is closely related to the timing of the EITC arrival. EITC recipients tend to increase their household spending, especially on durable goods (Barrow and McGranahan, 2001; Adams et al., 2009), consume more healthy food (McGranahan and Schanzenbach, 2014) and use more healthcare services (Hoynes et al., 2014; Niedzwiecki, 2013) in the February when they receive the credit.

[10]One recent paper finds that the tax refund provides liquidity for EITC-eligible job losers at the beginning of their unemployment spell (LaLumia, 2013). The author exploits the timing of the EITC refund to examine whether the unemployment duration (i.e. job search intensity) of EITC recipients is sensitive to the provision of a lump-sum transfer at the beginning of their unemployment spell and then estimates the liquidity effect of unemployment insurance. She finds that EITC recipients who become unemployed in February have longer unemployment duration than those entering unemployment in other months of the year.

ment group is based on family characteristics in the previous year. It is predetermined and cannot change during the year. My results show that the receipt of EITC causes the labor supply of married women to show a sharp drop in February and a substantial decline in January, March and April. This pattern is largely coincident with the payment timing of the EITC. In contrast, the receipt of the EITC has little impact on the timing of labor supply for married men and single women.

Second, I restrict my sample to those receiving EITC payments and utilize the variation in the amount they receive in a given month, which is predetermined by the time the labor supply decision is made, so as to quantify the impact of the EITC receipt on recipients' monthly labor supply. My estimates indicate that receiving a $1,000 EITC payment significantly reduces a married woman's likelihood of working, by 1.6 percentage points, in the month of the EITC arrival, from a base of 47%.[11] The income elasticity of labor supply for married women based on this estimate is around -0.06, which lies between the estimates in the previous studies (Blau and Kahn, 2007; Heim, 2007). In line with the results from the triple differences estimation, I do not find any statistically detectable effect of receiving a $1,000 EITC payment on the probability of working for either married men or single women.

I conduct several subgroup analyses to explore possible causes of my results. First, I investigate why there are very different labor supply responses to the receipt of EITC across married women, married men and single women. The subgroup analysis reveals such patterns to be due to the fact that the majority of married women are secondary earners in their families and not to gender differences in labor supply patterns among EITC recipients.[12] I find that married women who are secondary earners significantly reduce their labor supply in response to the receipt of EITC but those who are primary earners do not. Interestingly, a similar pattern also

[11] I use the likelihood of working in October, a month in which little of the EITC is disbursed, to represent the baseline mean for the treatment group.

[12] The definition of a secondary earner is an individual who earned less income than her spouse in the previous year. Single women are primary earners.

emerges in the sample of married men. These results suggest that families might adjust the labor supply of secondary earners in response to receiving an anticipated EITC payment. This result is consistent with findings in previous studies (Cullen and Gruber, 2000; Kohara, 2010) suggesting that the labor force participation of secondary earners is sensitive to changes in family resources.[13]

Next, I analyze why the family labor supply changes at the time of receipt of the anticipated EITC payment. My results suggest that it could be due to the presence of liquidity constraints and myopia among EITC-eligible families. The presence of liquidity constraints forces families to keep their labor supply high so as to maintain liquidity until receiving the EITC payment. Following previous studies, I conduct subgroup analysis by splitting the sample into those who are liquidity constrained and those who are less constrained. I find that married women from constrained families, such as families with low liquid assets or high mortgage-to-income ratios, exhibit a significantly negative labor supply response to the receipt of the EITC but those from less constrained families do not. However, liquidity constraints cannot fully explain why the receipt of the EITC causes married women to reduce their labor supply temporarily in February rather than smoothing out their labor supply in the months following the receipt of the cash. The observed pattern reveals that recipients could be somewhat myopic in planning for their future consumption.

Finally, I exploit month-to-month labor force transitions so as to understand the main cause of married women's decreased likelihood of working in February (relative to other months). Although the estimates are not precise, my results provide suggestive evidence that married women could temporarily leave their jobs without pay upon receiving the EITC refund in February.

This paper contributes to the existing literature in two ways. First,

[13]These phenomena have been addressed in field work. A respondent in Smeeding et al. (2000) vividly described her working status before receiving the EITC refund. As the authors explain: "She can pay off all her [back] bills, be caught up with all her bills and not feel stressed..... All she has to do is keep working until December. Then in January she can turn in her tax form so she can get that money."

it examines the labor supply response to the receipt of an anticipated cash transfer, which is largely unexplored in prior studies. Only two recent studies have studied the response of labor supply to an anticipated income change induced by cash transfers, both in developing countries where people have difficulty accessing credit. Edmonds (2006) finds that the timing of the anticipated public pension in South Africa is associated with the timing of child employment. Receipt of the pension reduces child labor supply and increases children's enrolment in school. Fernandez and Saldarriaga (2014) utilize exogenous variation in the time between the payment date of the conditional transfer program and the interview date of the household survey to find evidence of women's working hours in Peru declining upon receipt of a cash transfer. Due to data limitations, neither paper is able to conduct rigorous empirical investigations into the possible causes of their results. I use detailed asset and debt information in my data to investigate the mechanisms behind my findings. In addition, the present paper provides the first evidence on the causal effect of the receipt of a cash transfer on the timing of family labor supply in the context of a developed country.

Second, this paper represents the first attempt to analyze the short-run (i.e. intra-year) effects of the EITC on the labor supply. Since the EITC is a tax credit that subsidizes the earnings of low-income families, previous studies mainly focus on how a change in the level of the EITC payment affects the level of family labor supply.[14] They exploit various expansions of the EITC, which boosted the level of the EITC payment for eligible families between the mid-1980s and the mid-1990s, to evaluate the labor supply effect of the EITC. They find that EITC expansion resulted in an increase in the employment rate of single women(Eissa and Liebman, 1996; Meyer, 2010; Meyer and Rosenbaum, 2001) and a decline in the employment rate of married women (Eissa and Hoynes, 2004).[15] The present paper does not

[14]Theoretically, it should have both a substitution effect and an income effect on labor supply. The substitution effect comes about as a result of the EITC altering the family's labor supply decision by changing the marginal tax rate on earnings. The income effect results from the tax credit increasing the family's resources and leading individuals to decrease their labor supply (consume more leisure).

[15]Hotz and Scholz (2003) offer a literature review on behavioral responses to the EITC.

examine a change in the level of the annual EITC payment or labor supply. Rather, it focuses on whether the timing of EITC receipt affects the timing of family labor supply within any given year, keeping the annual EITC amount constant. My results clearly show that the timing of the disbursement of the EITC matters for the labor supply decision of married women.

The paper proceeds as follows. Section 2.2 briefly describes relevant features of the EITC. Section 2.3 discusses the data and sample selection process. Section 2.4 proposes the identification strategies. Section 2.5 presents my main results and robustness checks, and discusses possible mechanisms behind my findings. Section 2.6 concludes.

2.2 Background on The Earned Income Tax Credit

The EITC is a refundable tax credit for low-income working people, particularly those with children, in the US. In 1975, the EITC began as a small program but it has since grown into one of the largest anti-poverty programs in the US. In 2012, the US federal government spent $61 billion on the EITC, supporting more than 28 million families.

A taxpayer's eligibility for the EITC relies largely on her family's earned income (or adjusted gross income), number of qualifying children, and filing status during the tax year (i.e. previous year).[16] First, as the EITC is a policy tool aimed at encouraging the poor to work, a taxpayer must have positive earned income, defined as the sum of wage income and self-employment income. The final EITC payment depends on the minimum amount of credits based on either earned income or adjusted gross income (AGI).[17] Second, as with other means-tested transfer programs, a taxpayer's

Eissa and Hoynes (2006) provide a review focused particularly on the EITC's impact on labor supply.

[16] According to the Internal Revenue Service, "A tax year is an annual accounting period for keeping records and reporting income and expenses." Thus, a tax year usually refers to the previous year. From here on, I will use tax year and previous year interchangeably. There are three filing statuses: joint filing, single, and head of household. The first is for a married couple and the last two are for unmarried people.

[17] AGI is a taxpayer's total income from all sources, excluding non-taxable income such

AGI and earned income have to be below a particular income cutoff, which depends on the number of qualifying children and filing status. Third, a taxpayer with one or more qualifying child is eligible for a much larger amount of credit. Qualifying children must be under the age of 19 years, or 24 years if studying full-time, and must live with the taxpayer for at least half of the year. A small amount of credit is provided to childless taxpayers.

Figure 2.2 displays the EITC schedule for taxpayers with and without children.[18] The payment level is quite stable during my sample period of 1997 to 2012.[19] The EITC schedule consists of three regions: the phase-in region, where the tax credit increases at a given rate as earned income (or AGI) rises, the plateau region, where the tax credit stays constant at the maximum amount, and the phase-out region, where the tax credit declines at a specific rate for each extra dollar of income. The phase-in and phase-out rates depend on the number of qualifying children. For example, the phase-in rate for a taxpayer with one child is 0.34, so that one extra dollar of income would raise the EITC refund by 34 cents. The credit stops rising when it reaches the maximum amount and then stays unchanged until income hits the phase-out threshold. The credit will then start to phase out at the rate of 16 cents per dollar until it disappears entirely. Since 2002, married couples (i.e. married and filing jointly) have had a larger income threshold under which the maximum amount of credit can be given, which means that more tax credit is offered to married couples than to singles. In sum, there is a lot of heterogeneity in the amounts of credit paid to EITC recipients. For a taxpayer with two or more children, the maximum credit can account for 40% of family annual income. However, the maximum credit for a childless

as welfare benefits, minus any adjustments to income. The adjustments could be moving expenses, alimony paid, health savings account deductions, and so on.

[18]This is the EITC payment schedule for year 2007.

[19]In each year, the EITC payment is adjusted for inflation. The program was still being expanded somewhat during this period. For example, during from 2010 to 2013, as part of the American Recovery and Reinvestment Act, the EITC was temporarily expanded for families with three or more children. Therefore, the phase-in rate for families with three or more children became 45% of income (up from 40%). This change effectively raised the maximum credit for these families by around $600. The act also increased the income threshold at which credit begins to phase out for married couples to $5,000.

taxpayer only accounts for 5% of family annual income.

EITC payments usually arrive in the first quarter of the calendar year, mostly in February. This is because the EITC is part of the annual tax refund; EITC recipients receive their refunds in the first few weeks after filing their taxes, and the Internal Revenue Service (IRS) usually opens the filing window in mid-January.[20] This disbursement pattern is very different from those of other transfer programs and overpayment refunds, which tend to be distributed evenly over the calendar year.[21] Table 2.1 documents the share of total EITC disbursements that occur in each month, averaged across the years 1997 to 2012, based on various issues of Monthly Treasury Statements (MTS).[22] In each year, over 80% of EITC payments are disbursed between January and March. On average, the share of payments made in February is 56% and that in March is 22%.

Recipients could, during my sample period, obtain their payments even earlier than the MTS data show by using Refund Anticipation Loan (RALs), for which users were charged a very high fee (i.e. implicit interest rate) for the expedition of the receipt of their benefits. The service allowed a taxpayer to receive their refund immediately upon filing their tax return. Wu (2012) shows that around 18% of EITC recipients receive their tax refunds early via RALs. Moreover, according to McGranahan and Schanzenbach (2014), around 10% of EITC benefits were used to reduce the recipient's tax liability, and presumably such credits are received when the tax was paid. Taken together, these aspects imply that a substantial number of EITC recipients may have obtained their credits in January.

It was not necessary for a recipient to receive the EITC in the form of a one-time lump-sum payment during most of my sample period. Prior to

[20] In 2011, the window opened on January 14th.

[21] Other transfer programmes, such as Supplemental Security Income, Food Stamps, and Temporary Assistance to Needy Families, send the benefits out monthly. Individual income tax refunds are distributed evenly over March to May (Barrow and McGranahan, 2001).

[22] These are published by the Treasury Department's Financial Management Service. The information is available at `http://www.fms.treas.gov/mts/backissues.html`. As the IRS did not provide disbursement information in 1997, I used the 1998 distribution of disbursements to impute it.

2011, a recipient had the option of using "Advance EITC" to get back a portion of their expected credit each month over the calendar year prior to filing their taxes. However, this option was not the default and involved submitting paperwork to one's employer.[23] According to previous estimates (GAO, 1992), the take-up rate of Advance EITC was between just 0.5% and 3%. Given the evidence on liquidity constraints among EITC recipients, the extremely low participation in Advance EITC seems to have been a puzzle. Jones (2010) finds that the low take-up of Advance EITC did not result from a lack of information, administrative costs or stigma, and suggests that making Advance EITC the default option could have substantially increased its participation rate.[24] However, due to the very low usage rate, it has not been available since 2011.

2.3 Data and Sample

2.3.1 Data

The data I use come from the 1996, 2001, 2004, and 2008 panels of the Survey of Income and Program Participation (SIPP). The SIPP is a national representative survey of welfare program participation, employment and income dynamics, health insurance coverage, assets, liabilities, and related topics. The initial sample size for each panel is about 35,000 households and 100,000 individuals. Each panel is a longitudinal survey that follows the initially selected household members for at least three years and interviews them every four months.[25] In each interview, the respondent reports her or his labor participation and income sources for each of the preceding four months. Most of the information is reported at a monthly or quar-

[23] The maximum amount of advance credit that could be received was 60% of the maximum credit for a taxpayer with one child. The remaining credit was received after taxes were filed at the beginning of the next year. In 2009, a potential EITC recipient could obtain at most $1,826 through the Advance EITC.

[24] Based on my SIPP sample, I find that over 70% of EITC-eligible families hold at least one full-time jobs through out a year and 80% of them stay in the same jobs.

[25] Some panels, such as the 1996 panel, follow their sample for up to four years. In earlier years, the SIPP also had a short panel that followed selected household members for less than two years (e.g. the 1989 panel).

terly frequency. One exception is the variable indicating labor force status, which the SIPP provides weekly for each respondent. I use this information to construct my outcome variables.[26]

The SIPP data have two features making them especially suitable for this paper. First, the SIPP data have a longitudinal structure. This feature not only allows me to determine the treatment status for each person and precisely calculate the EITC payments by utilizing information on each family's income and number of qualifying children during the previous year, but also enables me to control for the unobservable time-invariant heterogeneity by including individual fixed effects in the regression.

Second, the SIPP also surveys household wealth and asset information once per year in its Assets, Liabilities, and Eligibility topical module.[27] The module provides the latest measurements of household assets, wealth, and debts as at the interview date, such as the value of deposits in bank accounts, stock and mutual fund holdings, home equity, vehicle equity, business equity, secured and unsecured debt, and mortgages. This information is particularly useful when I conduct subgroup analysis to explore possible explanations for my empirical findings by splitting the sample based on a family's tendency to be trapped by liquidity constraints. I use asset and wealth data from the topical module of the previous year to construct my measures of the family's liquid assets and mortgage-to-income ratio. This predetermined wealth information is used to form proxies indicating a family's liquidity situation in the current year.

2.3.2 Imputed EITC Payment

The SIPP does not provide valid information about the amount of EITC that each eligible family would have received.[28] I predict the amount of the EITC

[26] I will discuss my outcome variable in detail in Section 2.4.

[27] The waves of topical modules used in this paper are wave 3, wave 6, and wave 9 in the 1996, wave 3 and wave 6 in the 2001 and 2004 panels, and wave 4, wave 7 and wave 10 in the 2008 panel.

[28] SIPP indeed asks a question about the amount of EITC that a respondent receives in its tax topical module. However, the response rate of this question is fairly low, only 24%. In addition, some of answers to income questions in the tax module are inconsistent with

using information on family (earned) income, number of qualifying children, and filing (marital) status in the previous year. As mentioned before, the final amount of the EITC depends on the minimum amount of credits, based on either earned income or AGI. Since the SIPP has information about family earned income, I use this variable directly. However, the SIPP does not provide valid information about AGI. I use family income, the sum of earned income and unearned income, excluding non-taxable income, such as means-tested cash transfers,[29] to approximate AGI.[30]

Qualifying children must be under the age of 19 years, or 24 years if studying full-time, and must live with the taxpayer for at least half of the year. I use detailed information about the age of each family member, parent (father and mother) identifiers, school enrollment status, and number of months living with parents to calculate each family's number of qualifying children. According to EITC rules, married couples have to file their taxes jointly. Single individuals can choose either single or head of household filing status. Both filing statuses lead to the same EITC amount.[31] Thus, I use marital status to infer taxpayer's filing status when computing the EITC payment.

2.3.3 Sample

To improve the measurement of my outcome variable and the EITC payments for my estimation, I select my sample as follows. Table 2.2 displays the summary statistics of selected variables after each sample selection criterion has been applied. First, I require a respondent to have been followed in SIPP for at least two years. This criterion allows me to use the previous

those in the core SIPP files (Sisson and Short, 2001).

[29] The SIPP classifies family income sources into four categories: (1) earned income, (2) property income, (3) means-tested cash transfers, and (4) other income. The last three are unearned income but means-tested cash transfers are generally non-taxable income. Therefore, I define family income as the sum of earned income and unearned income, excluding means-tested cash transfers and use that to approximate AGI.

[30] Again, the SIPP has a question about the amount of AGI but the data quality is not good (i.e. low response rate and inaccurate numbers).

[31] However, filing as a head of household can provide more generous tax brackets and larger standard deductions than filing as a single.

year's information on family income and number of qualifying children to assign treatment status to an individual and infer the EITC amount that her family is likely to receive in the current year. For example, I use a respondent's 1996 information on family income and number of qualifying children to calculate the size of payment her family would have been likely to receive in 1997, and examine the impact of receiving the EITC payment on her intra-year labor supply pattern in 1997. Since I focus on low-income families, the estimated sample is restricted to those with positive earned income and family incomes below $40,000 in the previous year, which serves as the income cutoff for the comparison group (Column 1).[32] Following the previous EITC literature (Eissa and Hoynes, 2004; Eissa and Liebman, 1996), I conduct estimations separately for three subgroups that have been the populations of interest in previous studies: married women, married men, and single women (Column 2). Therefore, I only include those who are the reference person of their household or the spouse of the reference person. Note that a married couple filing their taxes together are from the same family and thus have the same predicted EITC amounts.

The basic unit interviewed in the SIPP is the household, and each household might have several families residing in it. In order to avoid the impact of the EITCs of other subfamilies within the same household, any individual living in a household with more than one family is dropped from the sample (Column 3). Furthermore, I restrict the sample to those aged from 20 to 55 so as to reduce the impact of retirement on my estimated labor supply responses (Column 4).[33]

Since my main focus is the intra-year change in labor supply, I require the sample to be observed for all 12 months in the years that I use for labor supply estimation (i.e. except in the first year of each panel) so as to mitigate concern about the impact of a change in sample composition on my estimates. This selection criterion ensures that my estimates identify

[32] I discuss the treatment and comparison groups in detail in Section 2.4.

[33] For married couples, this criterion is based on the wife's age. About 93% of the husbands are also within this age range. In Section 2.5, I report a robustness check using a sample in which both husband and wife are in the required age range (20 to 55).

changes in individuals' behavior instead of shifts in the composition of the sample (Column 5). Finally, for the first year of each panel, which is used only for determining the EITC amount and the treatment status in the following year, the sample can be observed for less than 12 months because a new SIPP panel might start after January. To obtain precise estimates of the EITC payments, I restrict the sample in these years to those with at least six months of observations (Column 6).[34]

The years I use for my estimations are 1997–1999, 2002–2003, 2005–2006, and 2009–2012. Note that I also use the first year of each panel (i.e. 1996, 2001, 2004, and 2008) to infer the treatment status and predicted EITC amount in the following years.[35] The final sample size comprises 25,564 individuals and 484,104 individual-month observations. From Table 2.2, one can see that the sample characteristics are fairly similar after each sample selection criterion is applied. The age restriction (Column 4) causes the biggest changes in the sample characteristics, causing the sample to have higher average earned income, a higher average number of children, a higher predicted EITC amount, a higher portion of EITC recipients, lower average wealth, and lower average liquid assets. According to statistics from the IRS, during my sample period, the average amount of the EITC payment was about \$1,974, which is quite close to the average value of the imputed credit amount, at \$2,130.[36]

2.4 Identification Strategy

In this section, I describe the empirical specifications used to examine the impact of the EITC receipt on the labor supply of low-income families. My identification strategy relies on the intra-year variation in the timing of EITC disbursement. EITC recipients receive their payments during the tax filing season, mostly in the month of February. I utilize this plausibly

[34]For those without twelve months of observations, I scale their incomes up to create an annual income. For example, for those with seven months of income information. I use the seven-month income multiplied by 12/7 to obtain an estimate of the annual income.

[35]These years are 1997, 2002, 2005, and 2009.

[36]All dollar amounts are in 2007 dollars.

exogenous timing of payments and the benefit rules of the EITC based on the previous year's information to estimate the causal effect of EITC receipt on low-income families' intra-year labor supply.

2.4.1 Triple Differences Estimation

I begin with triple differences analysis. This method compares the difference in labor supply for a treatment group, between February and the other months, to that for a comparison group, which is presumed to remove any shocks in February, other than the receipt of EITC payments, that might affect the labor supply decision of a treatment group. Following the prior literature (McGranahan and Schanzenbach, 2014; Niedzwiecki, 2013; LaLumia, 2013; Barrow and McGranahan, 2001), I define my treatment and comparison groups using predetermined information based on EITC benefit rules: (1) the family income in the previous year, and; (2) the number of qualifying children in the previous year. I estimate the following regression:

$$\begin{aligned} L_{imt} = \alpha + \beta_l LowInc_{it-1} + \beta_c Child_{it-1} + \beta_e EITC_{it} + \beta^{DDD} EITC_{it} \times Feb + \\ M + (LowInc_{it-1} \times M)\beta_{lm} + (Child_{it-1} \times M)\beta_{cm} + \delta_t + \nu_i + X_{imt}\psi + \varepsilon_{imt} \end{aligned} \tag{2.1}$$

where L_{imt} is my outcome of interest, the share of weeks worked by individual i in month m of year t. Since SIPP provides weekly labor force status,[37] I use the number of working weeks divided by number of weeks in a month to construct this variable: $L_{imt} = 1$ if individual i works for the full month; $L_{imt} = 0$ if individual i does not work at all during the month; $0 < L_{imt} < 1$ denotes cases in between.[38] The advantage of the above

[37]The SIPP questionnaire gives a respondent five choices for weekly labor force status: (1) with job or business, working; (2) with job or business, absent without pay; (3) with job or business, on layoff; (4) no job or business, looking for job or on layoff; (5) no job or business, neither looking for job nor on layoff. I use the first option to indicate that a respondent is working in a given week and the other four to indicate that she or he is not working.

[38]That is, an individual works for part of the month. For example, February has four weeks. If an individual only works for two weeks, I would assign $L_{imt} = 0.5$ to this

definition is that it can also capture changes in labor force status within a month. Later I report a robustness check of my estimates in which I used a different definition of the outcome variable.[39] The variable $LowInc_{it-1}$ refers to whether individual i's family income in year $t-1$ is greater than zero and less than the EITC income limit ($LowInc_{it-1} = 1$) or is greater than than EITC income limit and less than \$40,000 ($LowInc_{it-1} = 0$).[40] The income limit roughly corresponds to the maximum EITC-eligible income for the families with one child during my sample period. For a married couple, the income limit is \$36,000 and for single women it is \$33,000. The variable $Child_{it-1}$ refers to whether individual i has one or more qualifying children in year $t-1$ ($Child_{it-1} = 1$) or has no qualifying children in year $t-1$ ($Child_{it-1} = 0$).

The treatment group dummy $EITC_{it}$ can be expressed as an interaction term between $LowInc_{it-1}$ and $Child_{it-1}$. Therefore, $EITC_{it} = 1$ indicates that individual i belongs to the treatment group that is expected to receive high EITC payments in the year t, namely, those whose family income is below the EITC income limit and who have one or more children in year $t-1$. $EITC_{it} = 0$ denotes that individual i is in the comparison group that is expected to receive low EITC in year t due to either having too great an income or being childless in year $t-1$.[41] Note that the group assignment is based on the previous year's information, which is predetermined by the time an individual makes her labor supply decision. In other words, an individual's current labor supply cannot affect her treatment status.

observation.

[39] $L_{imt} = 1$ if individual i is working in any week during a month and $L_{imt} = 0$ otherwise.

[40] The income cut-off for the comparison group is chosen to narrow down the income difference between the two groups of families while retaining a sufficiently high sample size in the comparison group.

[41] I consider three comparison groups of individuals who are similar to those in the treatment group in many ways but receive very low EITC payments. The first consists of those individuals with a similar income level to those in the treatment group (i.e. individuals whose family income is below the EITC income limit) but no qualifying children. The second comprises those individuals with one or more qualifying child but whose family income in the previous year is just above the income limit and below \$40,000. The third comparison group includes childless individuals whose family income during the previous year is above the income limit but below \$40,000.

Figure 2.4 compares the distribution of EITC payments between the treatment and comparison groups for married couples and single women. One may notice that most individuals in the comparison group have predicted EITC payments of zero. On average, the predicted amount of EITC for married couples in the treatment group is about $2,450. However, those from the comparison groups only receive $140 on average. For single women, individuals from the treatment group are predicted to receive about $2,370 and those from the comparison groups just $70. Table 2.3 displays summary statistics of selected variables for the treatment group and comparison groups. As expected, the treatment group has larger family (earned) incomes, more children, and greater predicted EITC amounts than the comparison groups. In addition, the treatment group consists of more young, less wealthy, and less educated individuals than the comparison groups. However, except for the EITC-related variables, the differences in the covariates between the treatment and comparison groups are not statistically significant after controlling for individual fixed effects. Furthermore, I control for these covariates and individual fixed effects in all specifications, which substantially reduces the impact of these group differences on my estimates.

The variable Feb is a dummy for the month of February, when most EITC recipients receive their tax refunds. The key variable used for identification is $EITC_{it} \times Feb$, which indicates the February observation of individual i who is expected to receive a high EITC payment. Since the timing of the EITC payments is highly concentrated in February, the treatment group will experience a large cash influx in February due to the receipt of the EITC, which is assumed to be the only difference between the treatment and comparison groups during the year. Hence, I can attribute any February effect found in the treatment group to the impact of receiving the EITC. The coefficient of interest β^{DDD} represents the causal effect of the EITC receipt on the labor supply of individuals receiving high EITC payments in February.

Two assumptions are essential to ensure that β^{DDD} has a causal interpretation. First, in absence of EITC payments, the difference in labor supply between the treatment and comparison groups should be similar across all twelve months. In a later section, I conduct an event study analysis to investigate whether the differences in labor supply between treatment and comparison groups are similar across the months when very few EITC payments are disbursed. Second, the composition of the two groups cannot change across months. Since the membership of the groups is based on the previous year's information, there is no change in group composition within the current year. Moreover, the estimated sample is a fixed panel that follows the same individuals over twelve months.

I include a set of month dummies M so as to control for the monthly patterns in labor supply that is common to both treatment and comparison groups in all years, such as holiday-season jobs. The advantage of the triple differences regression is that it allows me to include more fixed effects that are related to the group-level seasonality in the labor supply. Since my treatment group consists of low-income individuals with children, the primary concern with my estimates is that the results could simply reveal monthly patterns in labor supply for specific groups, namely low-income individuals or individuals with children, regardless of the impact of receiving the EITC payment. Hence, I interact month dummies M with the low-income dummy $LowInc_{it-1}$ to further control for any monthly seasonality in labor supply that is specific to low-income individuals. Note that I use October as the baseline month since less than 1% of the total EITC disbursement is paid in this month. Similarly, to control for any monthly employment patterns for individuals with children, I also include group-specific month fixed effects for those who have qualifying children: $Child_{it-1} \times M$.[42] To control for common macroeconomic effects during my sample period, I include a series of year dummies δ_t. In addition, the panel structure of the data allows me to include individual fixed effects ν_i to control for any unobservable time-invariant differences in labor supply preferences between various individuals. Finally, to improve the precision of the estimates, I include a number of covariates X_{imt} that could affect an individual's labor supply: educational attainment, age, number of children below 18, family wealth, monthly state

[42] Again, I use October as the omitted month.

unemployment rate, state fixed effect, state-specific time trend, industry fixed effect, industry-specific time trend, a dummy denoting the interview month, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.[43]

The variable ε_{imt} represents an error term. Since I follow the same individuals over time, to account for possible serial correlation that might affect the estimation of the standard error, the standard errors in all regressions are clustered at the person level. All regressions are weighted using person-level weights provided by SIPP.[44]

2.4.2 Event Study Analysis

One possible concern in the above specifications is that I treat all months other than February as part of the comparison groups, i.e. as unaffected by the receipt of the EITC refund. However, nontrivial EITC payments are disbursed in other months, particularly January, March and April. Hence, the results from equation (2.1) might bias the estimates downward (in absolute value). To address this issue, I conduct an event study by replacing $EITC_{it} \times Feb$ with a full set of month effects M interacted with the treatment group dummy $EITC_{it}$ in regression (2.1).[45] The estimation is based on the following regression.

[43] The categories of an individual's educational attainment are high school drop-out, high school degree, and post-secondary education. Information about family wealth (2007 dollars) is taken from the Assets, Liabilities, and Eligibility topical module for the previous year. I use the information on a respondent's industry in the previous year and a quadratic time trend to construct an industry-specific time trend variable. I categorize individuals into five groups: agriculture, manufacturing, service, self-employed, and not working. The definition of a part-time worker is that the average weekly hours worked by the individual in the previous year were greater than zero but less than 20.

[44] In Section 2.5, I conduct robustness checks of my estimates by computing the standard error at different cluster levels and using unweighted regressions.

[45] Again, October is the omitted month.

$$L_{imt} = \alpha + \beta_l LowInc_{it-1} + \beta_c Child_{it-1} + \beta_e EITC_{it} + (EITC_{it} \times M)\beta_{em} + M + (LowInc_{it-1} \times M)\beta_{lm} + (Child_{it-1} \times M)\beta_{cm} + \delta_t + \nu_i + X_{imt}\psi + \varepsilon_{imt} \tag{2.2}$$

In practice, I plot the coefficients on the interactions between M and $EITC_{it}$ (October is the omitted month) to examine whether the monthly patterns of group differences in the labor supply are coincident with the timing of the EITC refund.

2.4.3 Individual Variation in Predicted EITC Payments

The triple differences approach has the virtue of having a source of identification that is quite transparent as it compares group-level outcomes. The drawback of this approach is that it compares differential treatment of relatively broad groups (i.e. high EITC versus low EITC) and assumes that treatment intensity (i.e. EITC payment amount) is the same within a group. However, there is substantial within-group variation in the amount of the expected EITC payment across individuals.

In this section, I utilize the EITC refund that the recipients are expected to receive in a given month to quantify the impact of the receipt of a $1,000 EITC refund on the recipient's labor supply during the month in which the refund is disbursed. To alleviate the concern over the comparability of labor supply behavior between EITC-eligible and EITC-ineligible individuals, I limit the sample to EITC recipients. The estimation is based on the following regression:

$$L_{imt} = \alpha + \beta^{IND} Refund_{it} \times Share_{mt} + \kappa_1 Refund_{it} + \kappa_2 Share_{mt} + M + X_{imt}\psi + \delta_t + \nu_i + \varepsilon_{imt} \tag{2.3}$$

In the spirit of Souleles (1999) and McGranahan and Schanzenbach

(2014), I construct a variable indicating the EITC payment that each recipient is predicted to receive in a given month in the following ways. First, the variable $Refund_{it}$ represents the EITC payment (in thousands of dollars) that individual i is predicted to receive in year t. Note that $Refund_{it}$ is predetermined as regards the dependent variable L_{imt} since the amount of the EITC payment is based on information from the previous year (i.e. year $t-1$). In other words, the current labor supply decision has no impact on the amount of EITC received. Second, since SIPP does not have information about when respondents receive their EITC payments, I use the aggregate-level measure of the share of annual EITC disbursement paid out in a given month m of year t, $Share_{mt}$, to approximate the date on which the recipient receives their EITC refund.[46] For example, $Share$ is set to 0.6 in February 2010 since 60% of the 2010 EITC refunds were disbursed in February. Using group-level refund timing instead of the exact dates on which individuals receive their EITC refunds could substantially reduce the endogeneity problem since the exact timing of the refund will largely depend on when an individual files her tax statement, which might be correlated with unobservable determinants of the individual's labor supply.

The key variable in this regression is the interaction term between $Refund_{it}$ and $Share_{mt}$, which represents the expected EITC payment in a given month for individual i. The coefficient of interest, β^{IND}, directly measures the effect of receiving a \$1,000 EITC payment on individual i's labor supply in the month in which the EITC payment arrives. This estimate is useful later, when I compute the income elasticity of labor supply based on this short-run change in labor supply induced by the EITC refund. Consistent with the triple differences analysis, I also control for month fixed effect M, year fixed effect δ_t, the individual fixed effect ν_i and the same set of covariates X_{imt} as before.

It has to be pointed out that most but not all of the tax refunds received by low-income families come from the EITC. Several previous studies (LaLumia, 2013; Romich and Weisner, 2000) show that the EITC could account

[46] The share data come from various issues of MTS.

for 70% to 80% of the tax refund for EITC-eligible families. Furthermore, the amount of the non-EITC refund, which comes from elements such as the child tax credit, may be positively correlated with the amount of the EITC. Therefore, it may be reasonable to assume that a larger EITC payment will be associated with a larger tax refund. The predicted amount of the EITC payment should provide a good approximation of the tax refund that low-income families will receive.

2.5 Results

2.5.1 Triple Differences Estimates

I start by presenting the estimates from the triple differences estimations. Table 2.4 reports the estimated coefficient on the key variable $EITC \times Feb$ in the triple differences estimation (equation (2.1)). Panels A to C present the results for married women, married men and single women, respectively. I begin by presenting the estimate from the basic triple differences regression controlling for the low-income group, the group with children, and month fixed effects as well as all possible two-way interactions between those three dimensions. Then, I gradually include the individual fixed effect, year fixed effect, state effect, and other individual characteristics that could determine the monthly labor supply,[47] so as to gain an understanding of the impact of adding other covariates to my estimates. The fact that the estimates do not change much across specifications with different sets of covariates is comforting, given the causal interpretation of the estimates.

In general, I find that the income seasonality induced by the receipt of the EITC refund leads to changes in the intra-year labor supply patterns of married women. My preferred specification (Column 5 in Panel A) indicates that, compared to married women who receive low amounts of credit, those who receive high EITC payments are, significantly, 2.9 percentage points less likely to work in the month of February than in other months. In sharp con-

[47]For a detailed list of the covariates in each specification, please see the note under Table 2.4.

trast, married men and single women who receive high EITC payments do not exhibit distinct likelihoods of working in February compared to the other months. The point estimates in Column 5, Panels B and C, suggest that the likelihood of working declines, in February (relative to other months), for married men and single women, by only 0.3 percentage points and 0.2 percentage points, respectively. Both estimates lack statistical significance.

2.5.2 Event Study Analysis

Next, I extend the triple difference estimation by replacing $EITC_{it} \times Feb$ with a full set of month effects M interacted with the treatment group dummy $EITC_{it}$ so as to examine whether the monthly pattern of group difference in labor supply largely follow the disbursement timing of the EITC. Figure 2.5 displays the coefficients on $EITC_{it} \times M$ and the corresponding 95% confidence intervals based on the sample of married women. One can see that the event study coefficients largely mirror the timing of the EITC disbursement. Compared to marred women who receive low EITC payment, the labor supply of those receiving high EITC payments drops much more in February and also shows a substantial decline in January, March and April (relative to October). Outside of these months, the difference in labor supply between the treatment and comparison groups is quite close to the baseline level in October. Figures 2.6 and 2.7 present the estimates for married men and single women, respectively. Consistent with the results from the triple differences estimation, no such pattern emerges for the married men and single women. Instead, the group differences in the outcome variable are quite similar over the twelve months.

2.5.3 Results from Individual Variation in EITC Payments

Finally, I report the results based on variation in the predicted size of the EITC payment for each recipient in a given month. This approach has the advantage of allowing variation in treatment intensity among EITC recipients. If the intra-year labor supply pattern of married women found in the previous section is driven by the receipt of the EITC, I should also find a more negative effect on the labor supply among those receiving larger EITC payments.

Table 2.5 reports the estimated coefficients on $Refund \times Share$ from equation (2.3). Again, I gradually include different sets of covariates so as to determine the impact of these covariates on my estimates. The estimates across the specifications are fairly independent of the introduction of different covariates. My preferred estimates (Column 5) suggest that the receipt of a \$1,000 EITC payment reduces the proportion of married women working by 1.6 percentage points in the month in which the EITC is received. Since the baseline mean of the outcome variable is 47%, the estimated decrease represents a 3.4% decline in the mean.[48] In line with my triple differences results, receiving a \$1,000 EITC payment does not have a statistically detectable impact on the share of weeks worked by either married men or single women in the month when the EITC is paid out.

2.5.4 Robustness Checks

In this section, I examine the sensitivity of my result to a variety of alternative sample selection criteria and empirical specifications. Table 2.6 displays several of the resulting estimates. The first row presents the estimates based on triple differences regression and the second row the estimates that utilize individual variation in predicted EITC payments. Column 1 presents the results for a sample with a lower age cut-off of 50. This sample selection further alleviates concerns over the impact of retirement on labor supply. In both specifications, the results suggest that this change has little impact on the estimates.

Next, I address the fact that the baseline sample is restricted to married women aged 20 to 55 while their spouses might not be between the ages of 20 and 55. In fact, 7% of the married women in the baseline sample had spouses

[48] Using hours of work and rates of pay, I can provide some insights into the extent of labor income offset by the EITC payments. The average hours of work for married women is \$105 hours per month and the average rate of pay is \$12. Therefore, my estimates in Column 5 of Table 2.5 implies \$1000 EITC payment can offset monthly earned income by \$20. The income replacement rate based on this calculation is around 2%.

aged above 55. Column 2 presents results based on a sample excluding married women whose spouses are outside of the specified age range. The estimates are quite similar to my baseline cases. Column 3 shows that the estimated coefficients from an unweighted regression are smaller than my main estimates (in terms of absolute value), although the point estimates are statistically indistinguishable from my main estimates. In Column 4, I redefine my outcome variable as follows: $L_{imt} = 1$ for an individual who works in any week during the given month and $L_{imt} = 0$ otherwise. This definition is more comparable to the outcome variable of labor supply in previous studies that use annual data but it ignores within-month variation in labor supply. Again, this change has little impact on my estimates.

Column 5 of Table 2.6 presents statistical inferences based on a different clustering level of standard errors. Since the policy variation I use is at the group-month level,[49] I present the standard errors clustered on the group-month cells to account for any dependence of the unobservable error within the group-month level.

There is a potential concern that my estimates could be confounded with fluctuations in labor demand due to holiday season jobs. Workers in these jobs are usually hired in the fourth quarter of the calender year and might quit their jobs in the first quarter of the following year. To alleviate this concern, Column 6 of Table 2.6 presents estimates based on a sample that excludes those who worked in the retail industry at the end of the previous year. This restriction reduces my sample by around 8% but has little impact on my estimates.

Some states have supplemental state EITCs. The size of state EITCs vary across states and time, which generates an additional source of variation in EITC payments. The last column of Table 2.6 presents the estimate based on regression 2.3 incorporating this state-level variation. I find that this estimate is quite similar to my baseline estimate.

[49] I have four groups. One is the treatment group and the other three are comparison groups. Therefore, the total number of group-month cells is 48 (4 groups x 12 months).

2.5.5 Discussion: Magnitude of the Estimates

In this section, I begin by discussing the estimates obtained from the two empirical approaches and then compare their magnitudes to estimates from the prior literature. Specification (2.3) suggests that, on average, the receipt of a \$1,000 EITC payment leads to a reduction in the likelihood of married women working by 1.6 percentage points during the month in which the EITC is received. The magnitude based on this specification indeed provides a similar qualitative conclusion to my triple-differences estimation. I present a simple calculation to confirm their similarity. First, the triple-differences estimation (specification (2.1)) suggests that the receipt of the EITC refund leads married women from the treatment group to be 2.9 percentage points less likely to work in February than in other months. Additionally, the average gap in EITC payments between the treatment and comparison groups in the triple differences estimation is around \$2,310. Note that around 56% of the EITC is paid out in February. Therefore, the estimate based on triple differences regression implies that receiving a \$1,000 EITC refund could reduce the proportion of married women working by 2.2 percentage points.[50] The estimates based on these two approaches are fairly close.

One way to think about the magnitude of my estimates is to calculate the income elasticity of labor supply and then compare it to the estimates reported in previous studies. The unearned income of married women is computed using the secondary earner assumption. That is, it is equal to the husband's monthly earned income plus the family's monthly unearned income. Since the receipt of the EITC has little impact on married men's labor supply, it could be reasonable to assume that the average size of monthly unearned income is unrelated to the EITC refund. The mean value of monthly unearned income for married women is around \$1,847. My estimate suggests that receiving a \$1,000 EITC refund could significantly reduce the proportion of married women working in the month in which the refund is received, by 1.6 percentage points from the base of 47%. In other words, a 54% increase in unearned income could lead to a 3.4% decline in the like-

[50] This is derived from a simple calculation: $\frac{0.029}{2.31 \times 0.56} = 0.022$.

lihood of working in the month of payment arrival. This implies that the income elasticity of labor supply for married women is around -0.06.

My estimated income elasticity is largely consistent with the findings in previous studies. McClelland and Mok (2012) provide an up-to-date review of labor supply elasticities. They point out that the previously estimated income elasticities of employment among married men and single women tend to be quite small, namely, close to zero. However, the responsiveness of the employment of married women to income changes is substantially larger than that for married men and single women. Heim (2007) finds the income elasticity of employment for married women to be between -0.13 and -0.05. Blau and Kahn (2007) estimate the income elasticity for married women to be about -0.1. Both studies rely on cross-sectional variation in unearned income. A few recent studies use more exogenous variation in income from randomized experiments, such as lotteries, to get more credible estimates of the income elasticity. Jacob and Ludwig (2012) use a randomized lottery for housing vouchers and estimate the income elasticity among lower-income individuals who apply for housing assistance to be -0.09.

One caveat should be noted when comparing my results to those in the previous literature. My estimated elasticity relies on a higher-frequency change in income and labor supply than previous studies have done. I exploit the *monthly* change in income induced by the tax refund and study the impact of this short-run income change on an individual's *monthly* working decision. However, most prior studies have utilized *annual* changes in income and labor supply to estimate income elasticity, meaning that their estimates could represent the relatively long-run relationship between income and labor supply. With this caveat in mind, my estimates are generally similar in magnitude to the previously estimated income elasticities of labor supply.

2.5.6 Mechanisms behind the Findings

Secondary Earner

I examine why married women's labor supply responds to the receipt of an EITC refund but married men's and single women's do not. One possible explanation is most married women are the secondary worker in a family. Prior studies find that labor supply of a secondary earner is quite sensitive to changes in family resources (Cullen and Gruber, 2000). The "added worker effect" hypothesis holds that, under an imperfect credit market, the secondary earners, typically married women, in families could provide transitory earning sources to smooth out household spending whenever families face temporary shortages of liquidity (e.g. if the family has a mortgage commitment). Secondary earners may then exit the labor market once the family's need for liquidity is met (Goux and Petrongolo, 2014; Heckman and MaCurdy, 1980; Kohara, 2010; Lundberg, 1980; Mincer, 1962).[51] Hence, one should expect a secondary earner to exhibit a more negative labor supply response to the receipt of the EITC than a primary earner. Another potential explanation for my findings is that female EITC recipients have specific intra-year patterns in labor supply that are coincident with the timing of the EITC disbursement.

I use information on individual earnings in the previous year to define the primary and secondary earners within each family. An individual who had lower annual earnings than her or his spouse during the previous year is classified as the secondary earner. I begin by focusing on married couples and estimate specification (2.3) for the following four subgroups: married women who are primary earners, married women who are secondary earners, married men who are primary earners, and married men who are secondary earners.

The first four columns in Table 2.8 display the coefficients on $Refund \times Share$ for the above four subgroups. The estimates in Columns 1 and 2 suggest that the negative labor supply response to EITC receipt for married women found in the previous section is exclusively driven by those who are secondary earners in their families. On average, upon receiving a $1,000 EITC payment, married women who are secondary earners are significantly less likely to work, by 1.7 percentage points in the month in which the EITC

[51] For example, Cullen and Gruber (2000) find that more generous unemployment insurance would "crowd out" the labor supply of married women who face a temporary reduction in household resources due to the unemployment of their husbands.

refund is received. In contrast, those who are primary earners only show an insignificant decreased likelihood of working, of 0.7 percentage points. Interestingly, a similar pattern arises in the sample of married men. A married man who is his family's secondary earner exhibits a reduction of 1.6 percentage points in his working likelihood in the month in which he receives a \$1,000 EITC refund. The magnitude of this estimate is fairly close to the estimate for married women but is not statistically significant due to the small sample size for this group. The above results imply that the negative labor supply response to the receipt of the EITC for married women could result from the different gender roles due to the division of labor within families, rather than gender differences in intra-year labor supply patterns among EITC recipients.

To further compare the two possible channels that could underlie my results, I pool the whole sample together (including single women)[52]and "horse race" the "secondary earner" channel against the "gender difference" channel by interacting the intercept, a set of month dummies, and the predicted monthly EITC amount with an indicator for being female ($Female$) and a dummy indicating that a person is a secondary earner ($Second$), respectively, in specification (2.3). The first row in the last column of Table 2.8 suggests that receiving a \$1,000 EITC payment leads the baseline group's likelihood of working in the month when the EITC is received to decline by 0.3 percentage points, insignificantly.[53] The "secondary earner" channel coefficient (in the second row) reveals that those who are secondary earners in their families will be an additional 1.4 percentage points less likely to work in a month in which they receive a \$1,000 EITC refund. The point estimate is significant, with a p-value of 0.03. However, the "gender difference" channel coefficient (in the third row) suggests that there is no statistically detectable additional impact of being female on the probability of working in the month in which the refund arrives, after controlling for the effect of being a "secondary earner."

[52]By definition, all single women are the primary earners in their families.

[53]The baseline group consists of married men who are primary earners.

Liquidity Constraints

Next, I analyze why families reduce labor supply upon receipt of an anticipated EITC payment. The leading explanation for the observed behavior is the presence of liquidity constraints preventing families from borrowing future income to finance current spending. In this situation, families may keep their level of labor supply high to improve their liquidity until their tight budget is loosened by the receipt of the EITC refund. If liquidity constraints do play an important role in determining the family labor supply, one would expect the negative labor supply response of married women to the receipt of the EITC to be driven by those women from "more constrained" families.

I use two proxy variables to indicate the family's tendency to be liquidity constrained: liquid assets (i.e. the value of bank deposits in the previous year) and the mortgage-to-income ratio (i.e. the amount of the mortgage divided by total family income in the previous year). Both variables are computed at the family level for the calendar year before EITC receipt. Following the standard methodology in the prior literature (Parker, 2014; Johnson et al., 2013; Johnson et al., 2006; Souleles, 1999; Zeldes, 1989), for each variable, I divide the sample into two sets of individuals: those likely to be liquidity-constrained and those likely not to be. I use *Constrained* to denote membership of the liquidity-constrained group and interact it with the intercept and the expected EITC amount in a given month in specification (2.3). Hence, the additional labor supply response to the receipt of a \$1,000 EITC payment for liquidity-constrained individuals would be identified by the interaction between the indicator for the constrained group and the predicted monthly EITC payment.

Individuals with low liquid assets could be unable to draw down their wealth to smooth out their spending. In order to improve their family's liquidity, they are likely to adjust their labor supply, which could result in a greater negative response of labor supply to the receipt of the EITC payment. In the spirit of Parker (2014), I label those with liquid assets below the one-month average family income (i.e.\$2,000) as constrained families and the rest as unconstrained.[54] Column (1) of Table 2.7 shows how the labor supply response to an EITC payment varies according to the liquid assets held. The first row indicates that the receipt of a $1,000 EITC payment reduces the likelihood of working for married women with high liquid assets by 0.3 percentage points in the month in which the EITC received and this estimate is not statistically significant. In sharp contrast, receiving a $1,000 EITC refund significantly lowers the proportion of married women with low liquid assets working by 1.7 percentage points (Row 3). This labor supply response is almost five times as large as that of married women from families whose liquid assets are above the average monthly family income (in absolute value). The point estimate of the group difference is statistically significant, with a p-value of 0.07 (Row 2).

Prior studies (Del Boca and Lusardi, 2003; Fortin, 1995) show that mortgage commitment is an important factor determining the labor force participation of married women. Furthermore, those who have large mortgages might also have limited borrowing ability since housing collateral is often used for borrowing (Mian et al., 2014; Mian and Sufi, 2011). I use the mortgage-to-income ratio to approximate the likelihood of families being bound by liquidity constraints. Families with high mortgage-to-income ratios may be under greater pressure to meet their mortgage commitment and have limited credit lines for borrowing additional money. Under these circumstances, married women with high mortgage-to-income ratios might enter the labor market temporarily to increase family liquidity, and their working decision may be sensitive to the change in family liquidity induced by the EITC refund. To investigate this hypothesis, the estimated sample is restricted to those who are house owners.[55] I classify families with a mortgage-to-income ratio of 1.5 or above as constrained and the remainder as less constrained. A mortgage-to-income ratio of 1.5 is around the median of the distribution of mortgage-to-income ratios. Column 2 of Table 2.7 shows how the labor supply response to an EITC payment varies according to the mortgage-to-income ratio. For married women with low mortgage-

[54] This value also divides the top 20% in the distribution of liquid assets from the rest.

[55] Around 56% of the EITC recipients in my sample are house owners.

to-income ratios, the receipt of a $1,000 EITC refund results in a decrease of 1 percentage point in the likelihood of working in the month in which the EITC is received (Row 1). However, the point estimate is not statistically significant. In contrast, the receipt of a $1,000 EITC refund significantly lowers the proportion of married women with a high mortgage-to-income ratio who work, by 2.7 percentage points (Row 3). The point estimate of the group difference is sizeable and statistically significant, with a *p*-value of 0.014 (Row 2).

Myopia

The above subgroup analysis suggests that the presence of liquidity constraints may be an important reason why married women reduce their labor supply when receiving anticipated EITC payments. However, the presence of liquidity constraints cannot fully explain why the receipt of the EITC causes labor supply of married women to have a temporary drop in February, and then revert back quickly to the normal level of labor supply. Assumed a household's utility function is concave so marginal utility of leisure is diminishing. Therefore, to maximize intertemproal utility, a household wants to keep marginal utility equal across time periods. In other words, if families are forward-looking but liquidity constrained, they should smooth out their labor supply (or leisure consumption) after receiving the cash (i.e. when they are no longer liquidity constrained). Thus, we should observe a small and persistent decrease rather than a large and temporary drop in labor supply following receipt of the EITC. This patterns reveals that the recipients could be somewhat present-biased and prefer consuming leisure at the time when receiving cash transfer rather than use this money for their future consumption.[56]

[56]There are several theories explaining such present-biased behavior. One possibility is that poor people could have hyperbolic discount rate (Angeletos et al., 2001), which results in time-inconsistent preferences. That is, current selves of EITC recipients want to enjoy leisure right after receiving payments but future selves of EITC recipients would like to keep their labor supply to finance future spending. Another explanation is that the presence of temptation goods causes people prefer to consume more today rather than save for tomorrow (Banerjee and Mullainathan, 2010). In contrast to normal goods, temptation

Month-to-Month Labor Force Transitions

Finally, I use detailed information about labor force status from the SIPP to investigate how the receipt of an EITC refund affects the recipient's month-to-month labor force transitions. This analysis will help us understand the main cause of the decreased likelihood of working for married women in February (relative to other months). In general, my results could be driven by the fact that the receipt of the EITC increases the likelihood of working-to-nonworking transitions or decreases the likelihood of nonworking-to-working transitions (e.g. those who work in January and then stop working in February or those who do not work in January and then keep not working in February due to the receipt of the EITC).

As mentioned before, SIPP classifies labor force status into five categories: (1) working, (2) temporary leave without pay, (3) temporary layoff without pay, (4) unemployment, and (5) out of the labor force. The last four categories are defined as nonworking. In addition, Figure 2.8 displays the total amount of income tax refunds paid out in each week from January to April, based on the Daily Treasury Statement, and clearly shows that a large amount of the income tax refund is disbursed in the 6th week and the 7th week, which corresponds to the second and third weeks of February. Since most EITC recipients obtain their refund in February, this weekly refund disbursement pattern is likely to reflect the timing of the receipt of the EITC refund. Therefore, it is quite possible that EITC recipients receive their credit in the third week of February.

Given the above information, I define the two outcome variables as follows and estimate specification (2.1): The first dependent variable, $P(W_{im} = 0|W_{im-1} = 1)$, is an indicator of whether individual i works in the third week of month $m-1$ (i.e. $W_{im-1} = 1$) and stops working in the third week in

goods, such as alcohol or cigarettes, generate utility only at the point of consumption and people do not value spending on tomorrow's temptation goods. For example, you may get pleasure from drinking alcohol today but you may think that it will be bad for your future selves to drink alcohol. Under the assumption that the proportion of temptation goods is decreasing as income increases, low-income people would prefer to consume today rather than tomorrow since a higher share of their future spending on temptation goods makes them unwilling to save for tomorrow consumption.

month m (i.e. $W_{im} = 0$), which measures the working-to-nonworking transition. Similarly, I measure the nonworking-to-working transition using the dependent variable, $P(W_{im} = 1|W_{im-1} = 0)$, an indicator of whether individual i does not work in the third week of month $m - 1$ (i.e. $W_{im-1} = 0$) but does work in the third week of month m (i.e. $W_{im} = 1$).

The first two columns in Table 2.9 report the estimated coefficients on $EITC \times Feb$ for the above two outcome variables. The estimated coefficient implies that the receipt of an EITC refund would increase the likelihood of a working-to-nonworking transition for married women by 1.24 percentage points in February compared to other months. On the other hand, the receipt of an EITC refund decreases the likelihood of a nonworking-to-working transition by 0.2 in February compared to other months. Most changes in labor force status are concentrated in the working-to-nonworking transitions, which is around six times as large as the nonworking-to-working transitions, although neither estimate is statistically significant. I further decompose the working-to-nonworking transition into more detailed changes in labor force status by estimating specification (2.1) with the following outcome variables, respectively: working to unpaid leave, working to temporary layoff, working to unemployment, and working to out of the labor force. When I do this, I find that most increases in working-to-nonworking transitions in February indeed come from working to unpaid leave transitions. The receipt of the EITC significantly increases the likelihood of working to unpaid leave transition for married women by 1.08 percentage points in February (relative to other months). In contrast, the likelihoods of other labor force transitions do not show significant differences between February and other months. Figures 2.9a to 2.9f, which plot the estimated coefficients on $EITC_{it} \times M$ from specification (2.2) (i.e. event study analysis) for each outcome variable, confirm my regression results. Although the estimates are not precise, my results provide suggestive evidence that married women could temporarily leave their jobs without pay upon receiving the EITC refund in February.

2.6 Conclusion

This paper utilizes the unique disbursement pattern and benefit rules of the EITC to examine the casual effect of the receipt of a cash transfer on the timing of family labor supply. My results show that income seasonality caused by EITC receipt leads to changes in the intra-year labor supply patterns of married women. On average, receiving a $1,000 EITC payment reduces the proportion of married women who work in the month of credit receipt by 1.6 percentage points from a baseline mean of 47%. The income elasticity of labor supply for married women based on these short-run changes in labor supply and income is around -0.06, which falls within the range of estimates in the previous literature that were obtained using longer horizons for employment and income changes. The analysis of month-to-month labor force transitions provides suggestive evidence that married women could temporarily leave their jobs without pay upon receiving the EITC refund in February. No such tax refund-induced intra-year labor supply emerges for married men or single women. The subgroup analysis suggests families might reduce the labor supply of secondary earners in response to receiving an anticipated EITC payment. In addition, my results suggest that the presence of both liquidity constraints and myopia among EITC recipients provide possible explanations for my findings.

Several interesting implications arise from my results. First, both this paper and previous studies consistently provide evidence of a liquidity constraint among those claiming the EITC. These results imply that providing more frequent payments prior to the tax filing year, such as through Advance EITC, should be an attractive option for low-income taxpayers and could substantially help liquidity-constrained recipients to smooth their spending and leisure throughout the year. However, the low participation rate in Advance EITC is still a puzzle in the literature, and it has not been an option since 2011. Recent studies (Jones, 2010) have made some progress toward solving this puzzle. In general, they find that the universally low take-up might not have resulted from recipients' lack of information about it, from the application process being too complicated, or from recipients' fear of stigma. Jones (2012) finds strong evidence on the presence of inertia among EITC recipients, suggesting making periodic EITC payments the default option could substantially encourage people to obtain their EITC throughout the year before tax filing. Future research on this issue is needed to aid the redesign of a feasible option for periodic EITC payments.

On the other hand, my results clearly show that married couples have more flexibility in terms of adjusting their labor supply so as to smooth out their spending than singles. One possibility for future research is to examine whether the response of household spending to EITC receipt (or the receipt of other anticipated income) varies by family structure, which would provide a more complete picture of how families smooth their consumption.

2.7 Figures

Figure 2.1: Real Federal Spending on EITC and TANF: 1975–2011

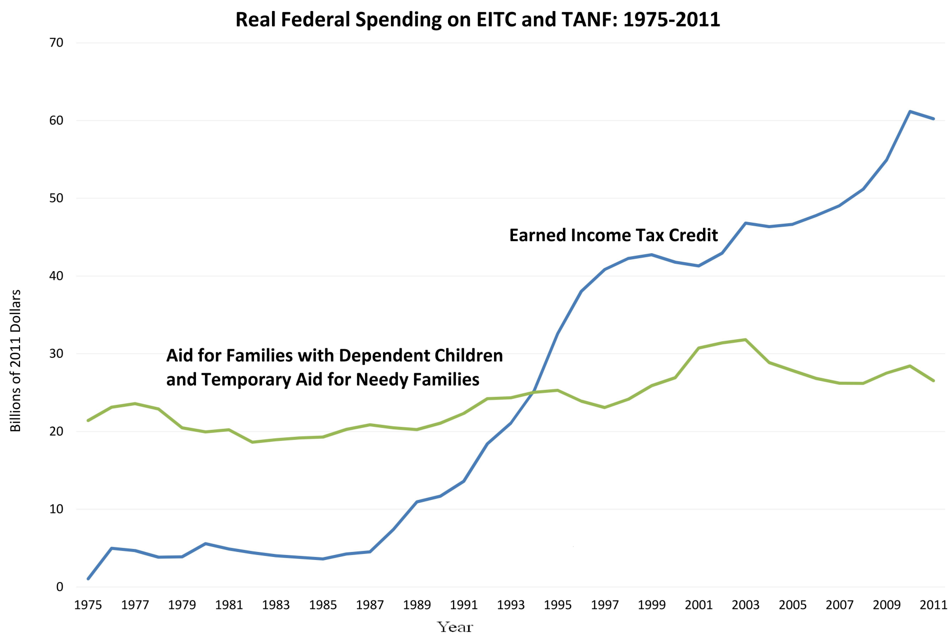

Notes: Data are from Tax Policy Center (2012)

Figure 2.2: EITC schedule (Tax Year 2007)

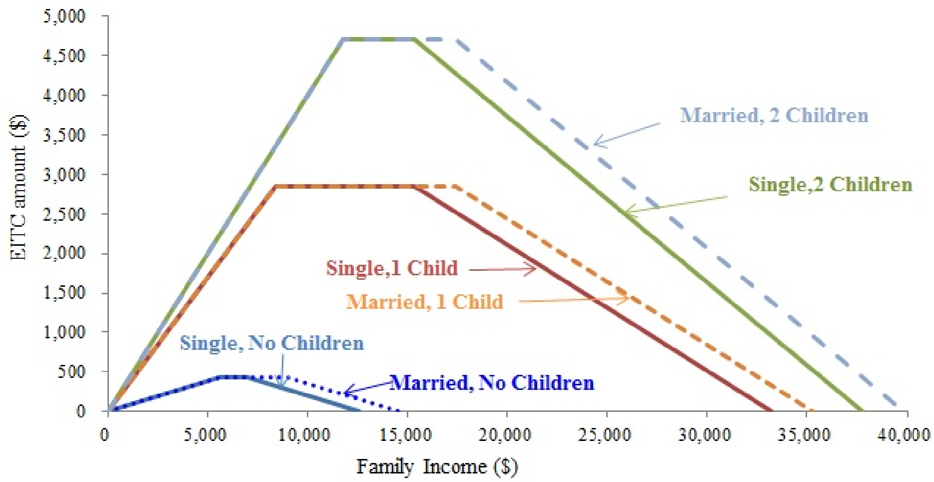

Notes: Data are from Tax Policy Center (2012). All dollar values are measured in 2007 dollars.

Figure 2.3: Share of Annual EITC Disbursements by Month

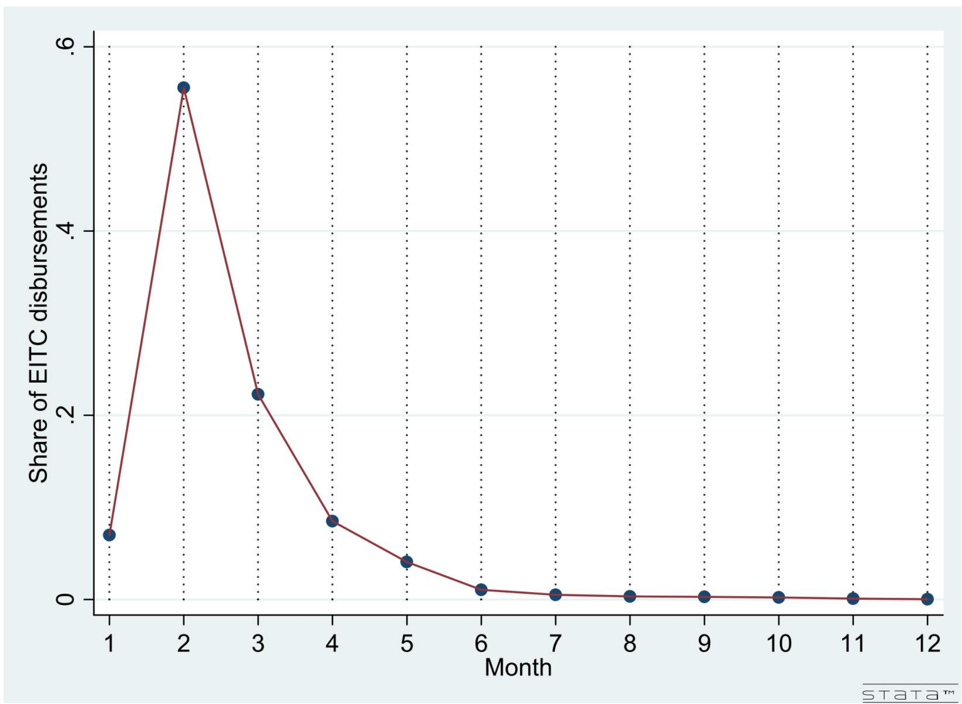

Notes: Data are from various issues of *Monthly Treasury Statements*. For each month and year, the fraction of the year's disbursements was first calculated. These fractions were then averaged by month across the years: 1997–1998, 2002–2003, 2005–2006, and 2009–2012. Because the IRS did not provide disbursement information in 1997, I used the 1998 distribution of disbursements to impute it.

Figure 2.4: EITC Amount by Treatment Status and Family Type

(a) Treatment Group: Married Couples

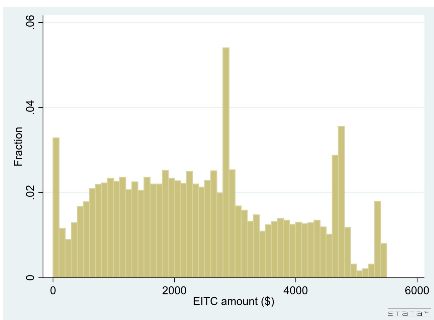

(b) Comparison Group: Married Couples

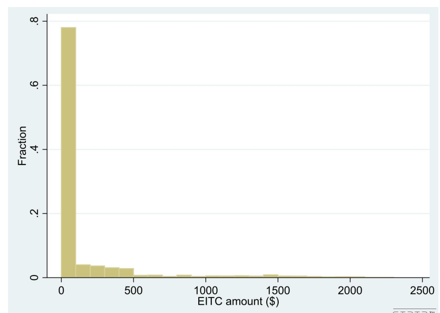

(c) Treatment Group: Single Women

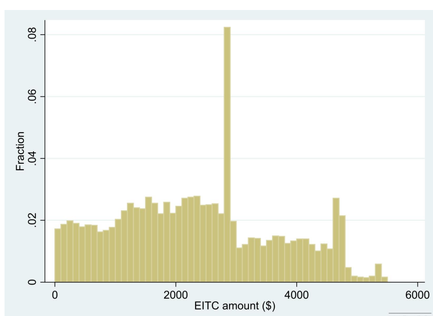

(d) Comparison Group: Single Women

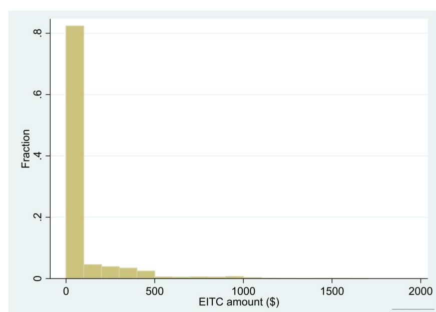

Notes: This table displays the distribution of predicted EITC amount by treatment status and family type. The horizontal axis indicates the predicted amount of the EITC. The vertical axis indicates the fraction of people within a specific income range. The bin width is $100.

Figure 2.5: The Impact of EITC on Intra-Year Labor Supply Patterns: Married Women

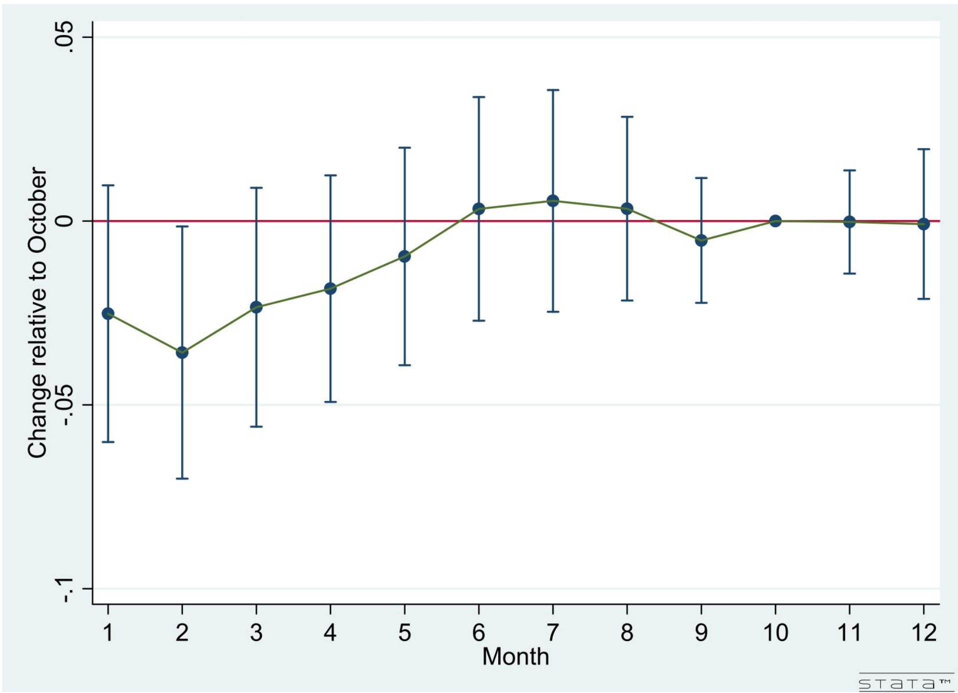

Notes: This figure shows coefficients on $EITC_{it} \times M$ and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to married women. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than $36,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than $36,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above $36,000 and below $40,000. (3) childless individuals that have incomes greater than $36,000 and below $40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below $36,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than $36,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Figure 2.6: The Impact of EITC on Intra-Year Labor Supply Patterns: Married Men

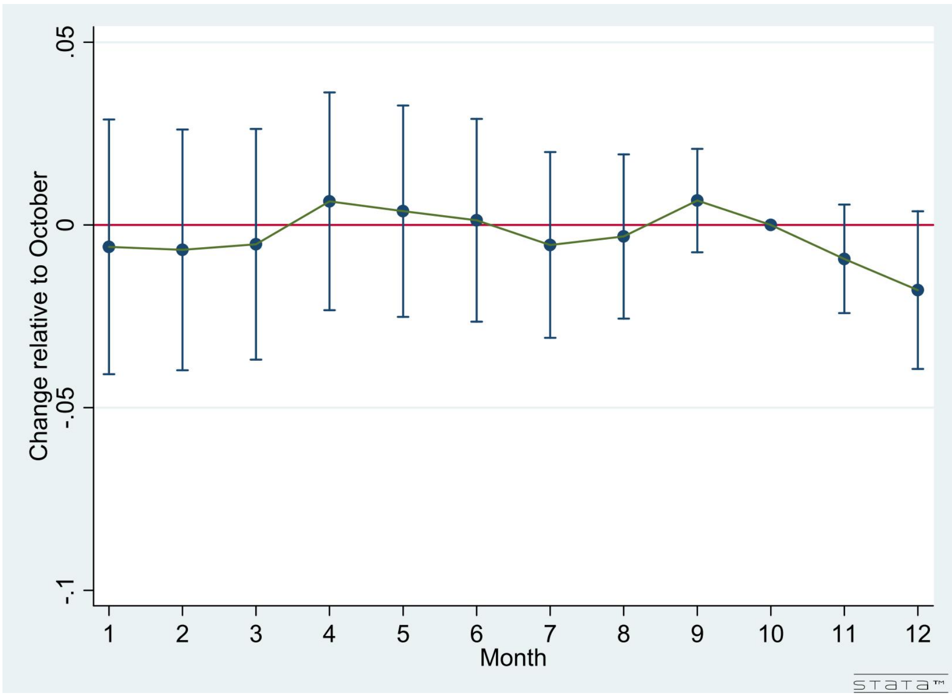

Notes: This figure shows coefficients on $EITC_{it} \times M$ and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to married men. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than \$36,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than \$36,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above \$36,000 and below \$40,000. (3) childless individuals that have incomes greater than \$36,000 and below \$40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below \$36,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than \$36,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Figure 2.7: The Impact of EITC on Intra-Year Labor Supply Patterns: Single Women

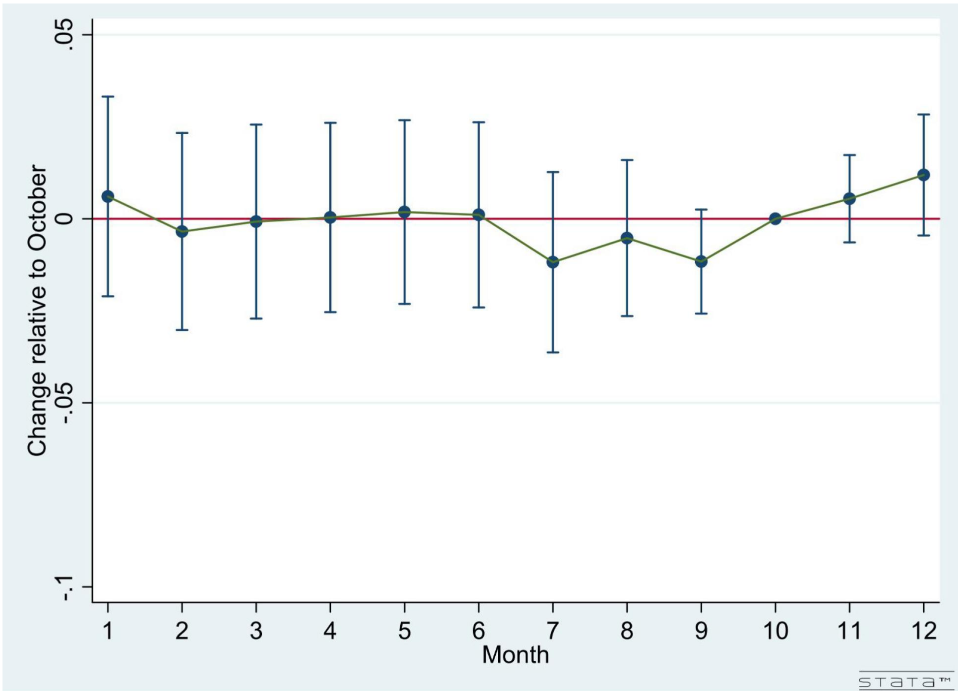

Notes: This figure shows coefficients on $EITC_{it} \times M$ and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to single women. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than \$33,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than \$33,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above \$33,000 and below \$40,000. (3) childless individuals that have incomes greater than \$33,000 and below \$40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below \$33,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than \$33,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Figure 2.8: Weekly Disbursement Patterns of Income Tax Refunds

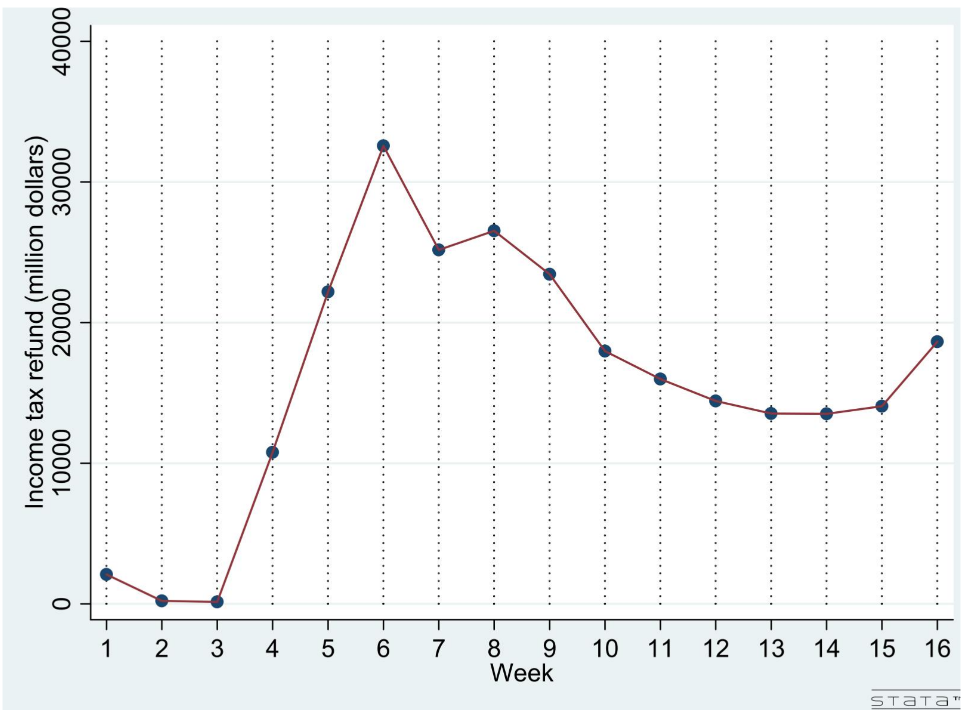

Notes: Data are from various issues of *Daily Treasury Statements*. The graph displays the average disbursement of income tax refunds during the first 16 weeks in a year. These amounts were averaged by week across the years: 1997–1998, 2002–2003, 2005–2006, and 2009–2012. Because the IRS did not provide disbursement information in 1997, I used the 1998 distribution of disbursements to impute it.

Figure 2.9: Month-to-Month Labor Force Transitions

(a) Working to Nonworking

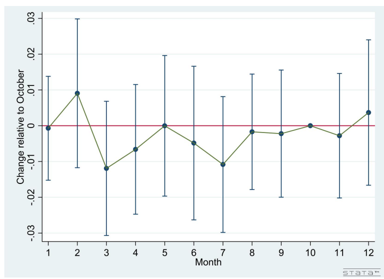

(b) Nonworking to Working

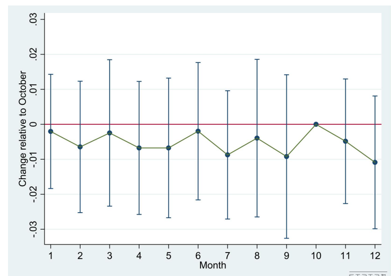

(c) Working to Unpaid Leave

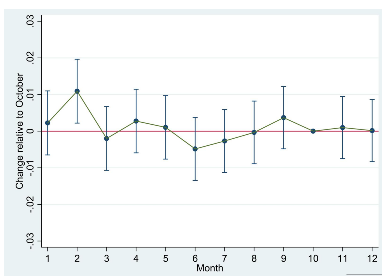

(d) Working to Temporary Layoff

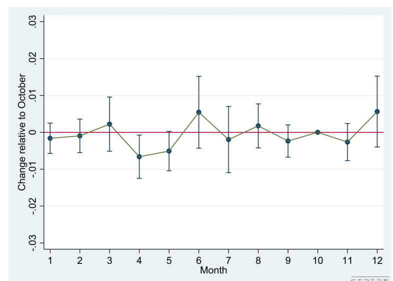

(e) Working to Unemployment

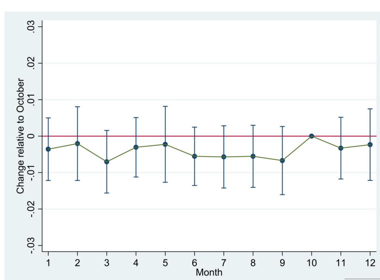

(f) Working to Out of Labor Force

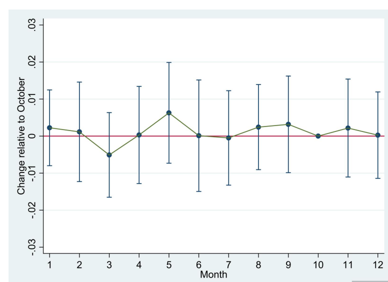

Notes: This figure shows coefficients on $EITC_{it} \times M$ and associated 95% confidence interval from specification 2.2 where the dependent variables are working to nonworking, nonworking to working, working to unpaid leave, working to temporary layoff, working to unemployment, and working to out of the labor force. The estimated sample is restricted to married women. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than $36,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than $36,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above $36,000 and below $40,000. (3) childless individuals that have incomes greater than $36,000 and below $40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below $33,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than $33,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

2.8 Tables

Table 2.1: Share of Annual EITC Disbursements by Month

	Percent of annual disbursements
January	7.0
February	55.5
March	22.3
April	8.5
May	4.1
June	1.1
July	0.5
August	0.3
September	0.3
October	0.2
November	0.1
December	0.0

Note: Data are from various issues of *Monthly Treasury Statements*. For each month and year, the fraction of the year's disbursements in that year was first calculated. These fractions were then averaged by month across the years: 1997–1999, 2002–2003, 2005–2006, and 2009–2012. Because the IRS did not provide disbursement information in 1997, I used the 1998 distribution of disbursements to impute it.

Table 2.2: Sample Selection: Summary Statistics

	(1) Below $40,000	(2) Reference person or spouse	(3) Only one family	(4) Age 20-55	(5) Observed 12 months	(6) First year information
Family income ($1,000)	24.17	24.53	24.82	24.45	24.90	24.92
	[10.13]	[10.04]	[9.98]	[10.14]	[10.02]	[10.03]
Earned income ($1,000)	19.78	20.07	20.14	22.14	22.50	22.51
	[11.02]	[11.10]	[11.17]	[10.67]	[10.58]	[10.59]
# of qualifying children	0.71	0.82	0.85	1.22	1.35	1.41
	[1.15]	[1.21]	[1.23]	[1.31]	[10.58]	[1.34]
EITC payment (> 0, $1,000)	1.74	1.83	1.86	1.97	2.10	2.12
	[1.49]	[1.49]	[1.50]	[1.48]	[1.48]	[1.48]
Working	0.67	0.68	0.67	0.72	0.71	0.71
	[0.46]	[0.46]	[0.45]	[0.44]	[0.44]	[0.44]
Wealth ($1,000)	93.50	93.81	97.52	66.34	71.31	70.79
	[876.92]	[927.03]	[930.28]	[627.81]	[802.97]	[816.96]
Liquid asset ($1,000)	4.80	4.77	4.97	2.85	2.84	2.81
	[21.54]	[21.16]	[21.93]	[15.36]	[15.19]	[14.89]
Amount of mortgage ($1,000)	25.53	25.10	25.24	27.19	29.91	29.61
	[56.48]	[55.31]	[55.31]	[57.59]	[60.02]	[59.45]
Age	43.16	44.18	45.78	39.00	39.52	39.50
	[15.38]	[14.52]	[14.92]	[9.98]	[9.76]	[9.75]
State Unemployment rate(%)	6.62	6.58	6.58	6.57	6.76	6.58
	[2.39]	[1.24]	[1.38]	[1.38]	[1.39]	[1.39]
EITC recipients	0.43	0.45	0.46	0.57	0.64	0.66
Female	0.56	0.57	0.58	0.56	0.66	0.66
Married	0.39	0.48	0.53	0.53	0.67	0.67
White	0.77	0.77	0.78	0.76	0.77	0.77
High school and below	0.72	0.72	0.72	0.70	0.71	0.71
Part time job	0.14	0.14	0.12	0.12	0.12	0.12
Secondary earner	0.20	0.24	0.25	0.34	0.34	0.34
Agriculture	0.06	0.06	0.07	0.07	0.06	0.06
Manufacturing	0.09	0.10	0.10	0.10	0.09	0.09
Service	0.57	0.56	0.58	0.58	0.57	0.57
Self-emplpoyed	0.10	0.10	0.10	0.10	0.10	0.10
Not working	0.18	0.18	0.15	0.15	0.18	0.18
# of EITC recipients	41,997	35,383	31,609	29,163	17,027	16,908
# of individual	91,607	72,438	64,240	49,580	26,185	25,564
# of individual-months	1,672,584	1,339,029	1,154,992	859,007	502,440	484,104

Note: SIPP data for years 1996–1999, 2001–2003, 2004–2006, and 2008–2012. Family income is the sum of earned income and unearned income, excluding the non-taxable mean-tested cash transfer. Family income, earned income, wealth, liquid asset, amount of mortgage, and EITC payment are in thousands of dollars. Family income, earned income, wealth, liquid asset, amount of mortgage, # of qualifying children, and EITC payment are based on the family level information in the previous year. All dollar amounts are in 2007 USD. Column (1) includes observations that have positive earned income and family income during the previous year below $40,000. In addition, they are followed for at least two years. Column (2) additionally requires sample to be married women, married men and single women. They are either reference person of the household or spouse of reference person. Column (3) additionally requires each sample to live in a household with only one family. Column (4) additionally imposes age restrictions: individuals with age 20 to 55. For a married couple, the age restriction is based on wife's age. Column (5) additionally requires individuals to be observed for all 12 months in the year that I use for estimating labor supply (i.e. except the first year of each panel). Column (6) additionally restricts the sample to those observed at least six months in the first year of each panel. Standard errors are reported in parentheses.

Table 2.3: Treatment and Comparison Groups: Summary Statistics

	Married Women		Married Men		Single Women	
	High-EITC	Low-EITC	High-EITC	Low-EITC	High-EITC	Low-EITC
Family income ($1,000)	23.97	30.49**	23.97	30.49**	17.55	25.11**
	[8.45]	[9.30]	[8.45]	[9.30]	[8.72]	[5.68]
Earned income ($1,000)	22.14	27.35**	22.14	27.35**	14.61	22.80**
	[8.94]	[10.89]	[8.94]	[10.89]	[8.62]	[10.59]
# of qualifying children	2.20	0.69**	2.20	0.69**	1.88	0.26**
	[1.13]	[1.20]	[1.13]	[1.20]	[1.02]	[10.59]
EITC payment ($1,000)	2.45	0.14**	2.45	0.14**	2.33	0.07**
	[1.44]	[0.37]	[1.44]	[0.37]	[1.32]	[1.59]
Working	0.44	0.58	0.82	0.77	0.77	0.87
	[0.49]	[0.48]	[0.49]	[0.48]	[0.40]	[9.93]
Wealth ($1,000)	70.97	112.92	70.97	112.92	24.17	46.60
	[207.1]	[1551.83]	[207.1]	[1551.83]	[90.68]	[816.96]
Age	36.19	41.02	39.13	44.10	37.11	41.09
	[7.96]	[10.06]	[8.94]	[10.06]	[8.24]	[9.75]
Unemployment rate(%)	6.68	6.61	6.68	6.61	6.39	6.46
	[2.39]	[1.24]	[2.39]	[1.24]	[1.39]	[1.39]
White	0.83	0.84	0.83	0.84	0.56	0.72
High school and below	0.77	0.71	0.77	0.72	0.72	0.55
Part time job	0.19	0.16	0.07	0.08	0.14	0.12
Agriculture	0.02	0.02	0.06	0.06	0.02	0.02
Manufacturing	0.06	0.06	0.26	0.21	0.08	0.08
Service	0.42	0.55	0.43	0.44	0.78	0.81
Self-employed	0.07	0.08	0.16	0.16	0.06	0.06
Not working	0.43	0.29	0.09	0.13	0.06	0.03
# of individual	5,202	4,063	5,202	4,063	3,984	4,770
# of individual-months	97,524	64,944	97,524	64,944	73,836	85,332

Note: SIPP data for years 1996–1999, 2001–2003, 2004–2006, and 2008–2012. Family income is the sum of earned income and unearned income, excluding the non-taxable mean-tested cash transfer. Family income, earned income, wealth, and EITC payment are in thousands of dollars. Family income, earned income, wealth, # of qualifying children, and EITC payment are based on the family level information in the previous year. All dollar amounts are in 2007 USD. The high-EITC group (treatment group) consists of those individuals that have one or more qualifying children and family income greater than zero and less than EITC income limit during the previous year. The low-EITC group (comparison group) comprise (1) those individuals that have family income greater than zero and less than EITC income limit during the previous year but have no qualifying child. (2) those individuals with one or more qualifying children but whose family income during the previous year is above EITC income limit and below $40,000. (3) childless individuals that have family income during the previous year is above EITC income limit and below $40,000. The income limit roughly corresponds to the maximum EITC-eligible income for the families with one child during my sample period. For a married couple, the income limit is $36,000 and for single women it is $33,000. Standard errors are reported in parentheses. Star indicates a significant difference across the preceding two columns after controlling individual fixed effects. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.4: Triple Differences Estimates

Dependent Variable:	Working				
	(1)	(2)	(3)	(4)	(5)
Panel A: Married Women					
EITC $\times Feb$	-0.0243**	-0.0292**	-0.0293**	-0.0295**	-0.0294**
	[0.0117]	[0.0116]	[0.0116]	[0.0116]	[0.0116]
R^2	0.025	0.786	0.787	0.790	0.793
Baseline mean			0.47		
# of individual			8,625		
# of individual-months			162,468		
Panel B: Married Men					
EITC $\times Feb$	-0.0046	-0.0035	-0.0035	-0.0037	-0.0031
	[0.0116]	[0.0117]	[0.0117]	[0.0117]	[0.0116]
R^2	0.018	0.705	0.705	0.708	0.711
Baseline mean			0.83		
# of individual			8,625		
# of individual-months			162,468		
Panel C: Single Women					
EITC $\times Feb$	-0.0022	-0.0042	-0.0042	-0.0038	-0.0021
	[0.0092]	[0.0092]	[0.0092]	[0.0091]	[0.0091]
R^2	0.022	0.654	0.654	0.657	0.658
Baseline mean			0.78		
# of individual			8,366		
# of individual-months			159,168		
Basic controls	√	√	√	√	√
Individual fixed effect		√	√	√	√
Year fixed effect			√	√	√
State effect				√	√
Other controls					√

Note: This table reports coefficients from triple differences regressions (equation (2.1)). The outcome variable L is share of weeks worked in a month defined as number of working weeks in a month divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The outcome variables regressed on the indicator for treatment group, as interacted with dummy for February. The treatment group consists of those individuals that have one or more qualifying children and family income greater than zero and less than EITC income limit during the previous year. The comparison group comprise (1) those individuals that have family income greater than zero and less than EITC income limit during the previous year but have no qualifying child. (2) those individuals with one or more qualifying children but whose family income during the previous year is above EITC income limit and below $40,000. (3) childless individuals that have family income during the previous year is above EITC income limit and below $40,000. The income limit roughly corresponds to the maximum EITC-eligible income for the families with one child during my sample period. For a married couple, the income limit is $36,000 and for single women it is $33,000. All dollar values are measured in 2007 USD. Column 1 control for treament group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below $36,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year greater than zero and less than $36,000, and month fixed effect. Column 2 additionally includes individual fixed effects. Column 3 additionally includes year fixed effects. Column 4 additionally includes state effects: state fixed effect, monthly state unemployment rate, state specific time trend (quadratic). Column 5 additionally includes other controls: educational attainment, age, number of children below 18, family wealth, industry fixed effects, industry-specific time trend, a dummy denoting the interview month, a dummy indicating that the individual worked part-time in the previous year, and month fixed effects specific to part-time workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.5: Estimates from Individual Variation in EITC Payments

Dependent Variable:	Working				
	(1)	(2)	(3)	(4)	(5)
Panel A: Married Women					
Refund $\times Share$	-0.0110**	-0.0144***	-0.0144***	-0.0161***	-0.0159***
	[0.0043]	[0.0043]	[0.0043]	[0.0043]	[0.0042]
R^2	0.021	0.777	0.778	0.782	0.785
Baseline mean			0.47		
# of individual			5,933		
# of individual-months			112,152		
Panel B: Married Men					
Refund $\times Share$	-0.0052	-0.0052	-0.0053	-0.0056	-0.0065
	[0.0043]	[0.0043]	[0.0043]	[0.0043]	[0.0042]
R^2	0.002	0.683	0.683	0.687	0.691
Baseline mean			0.83		
# of individual			5,933		
# of individual-months			112,152		
Panel C: Single Women					
Refund $\times Share$	0.0042	0.0053	0.0053	0.0056	0.00001
	[0.0054]	[0.0053]	[0.0053]	[0.0053]	[0.0053]
R^2	0.013	0.660	0.661	0.666	0.668
Baseline mean			0.78		
# of individual			5,079		
# of individual-months			92,688		
Basic controls	√	√	√	√	√
Individual fixed effect		√	√	√	√
Year fixed effect			√	√	√
State effect				√	√
Other controls					√

Note: This table reports coefficients from ordinary least squares regressions (equation (2.3)). The outcome variable L is monthly employment status defined as number of working weeks in a month divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The outcome variable is regressed on the imputed EITC amounts that an individual will receive $Refund$, as interacted with share of annual EITC disbursement paid out in a given month and year $Share$. The sample is restricted to EITC recipients. All dollar values are measured in 2007 USD. Column 1 controls for $Refund$, $Share$, and month fixed effects. Column 2 additionally includes individual fixed effects. Column 3 additionally includes year fixed effects. Column 4 additionally includes state effects: state fixed effects, monthly state unemployment rate, state specific time trend (quadratic). Column 5 additionally includes other controls: educational attainment, age, number of children below 18, family wealth, industry fixed effects, industry-specific time trend, a dummy denoting the interview month, a dummy indicating that the individual worked part-time in the previous year, and month fixed effects specific to part-time workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.6: Robustness Checks

Dependent Variable:	Working						
	(1) Age 20–50	(2) Husbands' age 20–55	(3) Unweighted regression	(4) Different dependent variable	(5) Stardard error group-month cluster	(6) No retail industry	(7) Add State EITC
Panel A: *Specification*(2.1)							
EITC $\times Feb$	-0.0279**	-0.0279**	-0.0246**	-0.0317***	-0.0294***	-0.0256**	
	[0.0140]	[0.0129]	[0.0106]	[0.0117]	[0.0021]	[0.0122]	
R^2	0.786	0.789	0.791	0.790	0.792	0.799	
# of individual	7,488	7,753	8,625	8,625	8,625	8,016	
# of individual-months	140,904	145,968	162,468	162,468	162,468	150,048	
Panel B: *Specification*(2.3)							
Refund $\times Share$	-0.0155***	-0.0163***	-0.0128***	-0.0160***	-0.0159***	-0.0149***	-0.0141***
	[0.0045]	[0.0044]	[0.0039]	[0.0043]	[0.0045]	[0.0043]	[0.0042]
R^2	0.782	0.783	0.783	0.781	0.785	0.783	0.792
# of individual	5,526	5,581	5,933	5,933	5,933	5,566	5,933
# of individual-months	104,340	105,612	112,152	112,152	112,152	104,484	112,152

Note: This table reports coefficients from ordinary least squares regressions (equation (2.1) and (2.3)). The outcome variable L is share of weeks worked in a month defined as number of working weeks in a month divided by total number of weeks in a month. Therefore, $L = 1$ if working for a full month, $L = 0$ if not working for a full month, and $0 < L < 1$ if working for a partial month. The outcome variable is regressed on the imputed EITC amounts that an individual will receive $Refund$, as interacted with share of annual EITC disbursement paid out in a given month and year $Share$. The sample is restricted to EITC recipients. Column 1 presents the results for a sample with a lower age cutoff of 50. Column 2 presents results based on a sample excluding married women whose spouses are outside of the age range (20 to 55). Column 3 shows the estimated coefficients from an unweighted regression. Column 4 shows that the estimated coefficients from an regression that uses different definition of outcome variable: $L = 1$ if working in any week during a month, and $L = 0$ otherwise. Column 5 shows the estimates using standard errors clustered on the group-month level (4 groups x 12 months). Column 6 presents results based on a sample excluding married women who worked in the retail industry in the end of previous year. Column 7 presents results that add variation in state EITC. In Panel A, all regressions control for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below \$36,000, month fixed effects for those who have qualifying children, month fixed effects for those who have family income during the previous year greater than \$36,000 and less than \$40,000, month fixed effects, individual fixed effects, year fixed effects, state fixed effects, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effects, industry specific time trend (quadratic), family wealth, a dummy indicating part time job workers in previous year, and month fixed effects specific to part time job workers. In Panel B, all regressions controls for $Refund$, $Share$, month fixed effects, individual fixed effects, year fixed effects, state fixed effects, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effects, industry specific time trend (quadratic), family wealth, a dummy indicating part time job workers in previous year, and month fixed effects specific to part time job workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.7: Subgroup Analysis Based on Tendency toward Being Liquidity Constrained (Married Women)

Dependent Variable:	Working	
	(1)	(2)
	Liquid Asset	Mortgage to Income
Refund $\times Share$	-0.0034	-0.0109
	[0.0043]	[0.0067]
Refund $\times Share \times Constrained$	-0.0135*	-0.0162**
	[0.0077]	[0.0066]
Row 1 + Row 2	-0.0169***	-0.0271***
	[0.0042]	[0.0066]
R^2	0.785	0.810
Mean of EITC (constrained)	$2,284	$2,140
Mean of EITC (less constrained)	$1,860	$1,981
# of individual	5,933	3,579
# of individual-months	112,152	62,579

Note: This table reports coefficients from ordinary least squares regressions (equation (2.3)). In addition, I use *Constrained* to denote membership in the liquidity-constrained group and interact it with the intercept and predicted EITC amount $Refund \times Share$. The sample is restricted to EITC recipients (Married Women). The outcome variable L is share of weeks worked in a month defined as number of working weeks in a month divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The outcome variable is regressed on the imputed EITC amounts that an individual will receive $Refund$, as interacted with share of annual EITC disbursement paid out in a given month and year $Share$. All dollar values are measured in 2007 dollars. All regressions control for $Constrained$, $Refund$, $Share$, month fixed effects, individual fixed effects, year fixed effects, state fixed effects, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effects, industry specific time trend (quadratic), family wealth, a dummy indicating part time job workers in previous year, and month fixed effects specific to part time job workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.8: Secondary Earner Channel v.s. Gender Difference Channel

Dependent Variable:	Working				
	Married Women		Married Men		Full Sample
	Primary Earner	Secondary Earner	Primary Earner	Secondary Earner	
Refund $\times Share$	-0.0067	-0.0171***	-0.0046	-0.0158	-0.0028
	[0.0094]	[0.0048]	[0.0046]	[0.0113]	[0.0042]
Refund $\times Share \times Second$					-0.0138**
					[0.0058]
Refund $\times Share \times Female$					0.0028
					[0.0056]
R^2	0.648	0.767	0.576	0.776	0.762
Baseline mean	0.85	0.36	0.90	0.54	0.67
# of individual	1,453	4,766	4,766	1,453	11,012
# of individual-months	23,520	88,632	88,632	23,508	316,980

Note: This table reports coefficients from ordinary least squares regressions (equation (2.3)). The outcome variable L is share of weeks worked in a month defined as number of working weeks in a month divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The outcome variable is regressed on the imputed EITC amounts that an individual will receive $Refund$, as interacted with share of annual EITC disbursement paid out in a given month and year $Share$. In the last column, based on equation (2.3), I also interact the intercept, predicted monthly EITC amount, and a set of month dummies with a indicator for female ($Female$) and a dummy indicating secondary earner ($Second$), respectively. The sample is restricted to EITC recipients. All regressions control for $Female$, $Second$, $Refund$, $Share$, month fixed effect, month fixed effect specific to female, month fixed effect specific to secondary earner, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating part time job workers in previous year, and month fixed effect specific to part time job workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 2.9: Month-to-Month Labor Force Transitions (Married Women)

	Month-to-Month Labor Force Transitions					
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Working to Nonworking	Nonworking to Working	Working to Unpaid Leave	Working to Temporary Layoff	Working to Unemployed	Working to out of Labor Force
EITC $\times Feb$	0.0124	-0.0021	0.0108*	-0.0003	0.002	0.00003
	[0.0086]	[0.0073]	[0.0056]	[0.0038]	[0.0036]	[0.0052]
R^2	0.071	0.069	0.074	0.056	0.065	0.062
Baseline mean	0.016	0.018	0.004	0.002	0.003	0.007
# of individual	8,625	8,625	8,625	8,625	8,625	8,625
# of individual-months	162,468	162,468	162,468	162,468	162,468	162,468

Note: This table reports coefficients from triple differences regressions (equation (2.1)). The sample is restricted to married women. The outcome variables across the columns are presented as follows: 1) an indicator for whether an individual work in the third week in the last month and stopped working in the third week in the current month; 2) an indicator for whether an individual did not work in the third week in the last month and started working in the third week in the current month; 3) an indicator for whether an individual worked in the third week in the last month and then took leave without pay in the third week in the current month; 4) an indicator for whether an individual worked in the third week in the last month and then had temporary layoff without pay in the third week in the current month; 5) an indicator for whether an individual worked in the third week in the last month and become unemployed in the third week in the current month; 6) an indicator for whether an individual work in the third week in the last month and moved out of the labor force in the third week in the current month; All regressions control for treament group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below $36,000, month fixed effects for those who have qualifying children, month fixed effects for those who have family income during the previous year greater than $36,000 and less than $40,000, month fixed effects, individual fixed effects, year fixed effects, state fixed effects, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating part time job workers in previous year, and month fixed effects specific to part time job workers. Standard errors are clustered at the person level and reported in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Chapter 3

Patient Cost-Sharing and Healthcare Utilization in Early Childhood: Evidence from a Regression Discontinuity Design

3.1 Introduction

Health conditions and medical treatments in early childhood are widely believed to have a substantial impact on health and labor outcomes in adulthood (Bharadwaj et al., 2013; Almond et al., 2011; Currie, 2009; Almond, 2006; Case et al., 2005; Currie and Madrian, 1999).[57] On the other hand, young children also bring about sizeable medical costs for their parents since they are vulnerable to diseases.[58] In line with this evidence, many public health insurance programs around the world subsidize healthcare service for young children by requiring relatively low patient cost sharing from this age group.[59] For example, the United States regulates the level of patient cost

[57] Several recent studies (Bharadwaj et al., 2013; Almond et al., 2011) present convincing evidence showing that early-life medical treatments can reduce mortality and even result in better long-run academic achievements in school. That is, health intervention in early childhood could be an investment with high returns.

[58] For example, in Taiwan, the number of outpatient visits for children under 3 years of age is around 20 per year. Compared with adults (12 visits per year), this age group has an especially high demand for healthcare service.

[59] That is, the share of healthcare costs paid out-of-pocket by the patient is lower.

sharing in Medicaid and the Children's Health Insurance Program (CHIP) to ensure that children from middle and low-income families can afford essential medical treatment.[60] Recently, due to tight budgets, many state governments have considered raising the level of patient cost sharing for Medicaid and CHIP, which has led to many debates on the possible impact.[61] Similarly, national health insurance in Japan and Korea offer children under 6 years of age a lower level of patient cost sharing than those above age 6, to promote health investments in early childhood.[62]

Low patient cost sharing has clear trade-offs. On one hand, low cost sharing can protect patients from financial risk induced by huge medical expenses. Better financial protection can help households smooth out their consumption and then increase the welfare of households. In addition, low cost sharing can make healthcare service affordable (i.e. income effect). Especially, low-income families might need healthcare service but cannot afford it because they may be liquidity constrained. On the other hand, low cost sharing could induce patients to overuse healthcare service. Since insured people do not pay the full cost of healthcare services, the optimal utilization of healthcare for an individual would be larger than the social optimum, leading to a loss of social welfare (i.e. moral hazard). To determine the appropriate level of cost sharing for children, we need to understand how cost sharing affects children's demand for healthcare service, namely, price elasticities of healthcare utilization.

[60] The federal requirement for Medicaid eligibility varies according to the children's age. For children under age 6 (young children), Medicaid eligibility requires family incomes to be lower than 133% of the federal poverty level (FPL). For children ages aged 6–19 (older children), family incomes is required to be below 100% of FPL. Thus, the coverage of Medicaid for children under 6 is much higher than for those above 6.

[61] Since the passing of the Deficit Reduction Act (DRA) of 2005, states have had the right to increase the level of cost sharing in public health insurance programs, such as Medicaid and CHIP, for specific populations and medical services (Selden et al., 2009).

[62] National health insurance in Japan covers almost all medical services, such as outpatient and inpatient care, for all citizens. The patient cost sharing for children under age 6 (pre-school age) is 20% of the original healthcare cost. For children above age six (school-age), patient cost sharing rises to 30% of medical costs. More details of Japanese national health insurance can be found at this web page:`http://www.shigakokuho.or.jp/kokuho_sys/kokuho_en.pdf`. In Korea, their national health insurance exempts cost sharing for inpatient services for children under age 6.

To date, very little is known about how young children's demand for healthcare services reacts to changes in the level of patient cost sharing. Previous studies mainly focus on price elasticities of adults' demand for healthcare (Cherkin et al., 1989; Selby et al., 1996; Rice and Matsuoka, 2004; Chandra et al., 2010a; Chandra et al., 2010b; Chandra et al., 2014; Shigeoka, 2014).[63] However, these estimates might not be valid for the demand for healthcare services of young children for two reasons.

First, the types of healthcare services used by adults and children are quite different. Children's outpatient visits are rarely for chronic diseases and mostly for acute diseases, which need timely treatment and should not be sensitive to a price change. In addition, the majority of children's inpatient admissions are for respiratory diseases, which can be treated with bed rest or medication. Previous studies have found the demand for this type of inpatient care is not price sensitive.[64]

Second, healthcare interventions in early childhood could substantially benefit an individual's later life, as addressed by recent studies (Bharadwaj et al., 2013; Almond et al., 2011). Given such high returns, parents might not be willing to adjust their children's medical care in response to price changes. Based on the above two reasons, we expect healthcare utilization for young children to be less price sensitive than that for an older demographic group.

Credible estimates of price elasticity for children still rely on evidence from the RAND Health Insurance Experiment (RAND HIE), which was an influential randomized social experiment conducted in the mid 1970s.[65]

[63] Shigeoka (2014) exploited the sharp reduction in patient cost sharing at age 70 in Japan and applied a regression discontinuity (RD) design to estimate the price elasticity of outpatient and inpatient visits by the elderly. He found the use of both health services to respond strongly to the price change with obvious drops at age 70. The estimated price elasticities were around -0.17 (outpatient) and -0.15 (inpatient). Chandra et al. (2014) used a cost sharing reform in Massachusetts as an exogenous variation in price and obtained a price elasticity of healthcare expenditure of around -0.15 for low-income adults.

[64] Shigeoka (2014) found that inpatient admissions for non-surgery were less price sensitive than those for surgery, especially elective surgery (e.g. cataract surgery). Also, he found that admissions for the respiratory diseases typically treated with bed rest or medication did not respond to a change in cost sharing at age 70 in Japan. Card et al. (2008) obtained similar findings for Medicare eligibility at age 65 in the United States.

[65] Before the passing of the DRA of 2005, state governments had little right to adjust the

Its sample comprised people of 62 years of age or under and randomly assigned participating households to different levels of patient cost-sharing (ranging from free care to 95% cost-sharing). The RAND HIE provided estimates of the price elasticity of healthcare utilization for children under 14 years of age (Leibowitz et al., 1985; Manning et al., 1981). It found that higher patient payments significantly reduced children's outpatient expenditure and utilization, but found mixed evidence of the cost sharing effect on children's demand for inpatient care.[66] The estimated price elasticity of the total healthcare expenditures was around -0.12.[67] However, the sample size for children in the RAND HIE was not big. Some estimates or subgroup analyses were not precise enough to confirm the presence or absence of a cost-sharing response (Leibowitz et al., 1985).[68] Additionally, the RAND HIE evidence is now over 30 years old. Both medical technology and the market structure have changed considerably during the past three decades. The varying healthcare environment could affect the way in which demand

level of patient cost sharing in their public insurance programs (i.e. Medicaid and CHIP) for children. Thus, there is little evidence on the effect of cost sharing on children's demand for healthcare service. To the best of our knowledge, only one recent study (Sen et al., 2012) has used the copayment change in the CHIP in Alabama, to analyze this issue. However, their study mainly relied on pre-/post-policy analysis, which suffers from the an estimation bias due to uncontrolled trends in children's medical utilization.

[66] For children under age 4, the RAND HIE found that inpatient care was price sensitive. Children assigned to a free plan had a significantly higher rate of inpatient admission than children assigned to 95% cost-sharing. For children aged between 5 and 13, no consistent pattern of a cost sharing effect on inpatient use was found (Leibowitz et al., 1985).

[67] The health insurance contracts in RAND HIE adopted non-linear pricing, which makes estimating price elasticity challenging. Specifically, the insurance plans required initial cost-sharing (free care, 25%, 50% and 95%) but had an annual stop-loss (Maximum Dollar Expenditure), in that the total out-of-pocket medical costs per year could not exceed 4,000 USD. Thus, the patient cost-sharing would fall to zero when annual out-of-pocket medical costs reached 4,000 USD. Such non-linear pricing imposes on patients different prices for the same health care at different times in the year. To summarize the estimated price elasticity, RAND researchers defined four kinds of price that patients respond to when making their healthcare decision: (1) the current "spot" price, (2) the expected end-of-year price, (3) the realized end-of-year price, and (4) the weighted-average of the price paid over a year (Aron-Dine et al., 2013). The price elasticity of children's healthcare mentioned here is calculated by defining price as definition (1).

[68] As Leibowitz et al. (1985) comment: "Because hospitalizations for children are infrequent, our estimates of hospital use have wide confidence intervals and we can be less certain than for outpatient care about the presence or absence of a cost sharing response."

for healthcare service changes in response to differences in price. Therefore, our paper fills this gap by providing the latest estimates of the price elasticity of children's healthcare utilization.

In this paper, we exploit a sharp increase in patient cost-sharing in Taiwan at the 3rd birthday that results from young children "aging out" of the cost-sharing subsidy. On average, turning age 3 causes a small increase in price per outpatient visit by about 60 NTD (or roughly 2 USD).[69] In addition, the increase in outpatient price at 3rd birthday is not uniform across different healthcare providers. Turning age 3 causes larger increase in outpatient price for teaching hospitals than clinic and community hospitals. This is because patients do not pay a copayment before the 3rd birthday. After the 3rd birthday, patients start to pay a copayment for each outpatient visit. Copayments for teaching hospitals are much larger than clinics and community hospitals. Finally, turning age 3 results in a much larger increase in price per inpatient admission from zero to 1,300 NTD (i.e. 40 USD). We use a regression discontinuity (RD) design to examine the causal effect of patient cost sharing on children's demand for healthcare service by comparing the expenditure and utilization of healthcare for children just before and after the 3rd birthday.

We obtain three key findings. First, a small increase in outpatient price at the 3rd birthday results in a sizeable reduction in outpatient utilization. The number of outpatient visits drops sharply at the 3rd birthday. The implied price elasticity of outpatient utilization is around -0.10. Second, the price increase at age 3 not only results in fewer outpatient visits (extensive margin) but also reduces the expenditure of each visit (intensive margin), namely, it induces patients to switch from high to low-quality providers (e.g. substitution of teaching hospitals with clinics or community hospitals). We find turning age 3 (i.e. paying copayment) reduces visits to teaching hospitals by 50% and most of the foregone visits are for less severe conditions.[70] Further investigating possible heterogeneous effects in detail, we also find

[69] 1 USD is equal to 32.5 NTD in 2006 prices.

[70] This result is due to copayments varying between health providers in Taiwan. We will discuss this issue in more detail in Sections 3.2 and 3.5.

preventive care and mental health services to have larger price responses than healthcare for acute respiratory diseases. Third, in sharp contrast to outpatient services, the demand for inpatient services does not respond to the price change at the 3rd birthday even if change in the inpatient price at age 3 is much larger than that in the outpatient price in terms of its level and percentage change. The estimated price elasticity of inpatient utilization is close to zero (about -0.004). This finding implies children's inpatient care could be quite necessary. Parents are unwilling to reduce a children's inpatient care even through they pay higher price after the children's 3rd birthday. The above findings suggest that the level of patient cost sharing for young children should differ depending on the healthcare service. For example, providing free inpatient care for young children does not stimulate excessive hospital use (i.e. moral hazard) but it might substantially reduce the financial risk for households. On the other hand, having a certain level of copayment for children's outpatient care is essential to avoid overuse of outpatient care, especially, at teaching hospitals.

This paper contributes to the research on patient cost sharing in three ways. First, the unique setting in Taiwan makes our estimates free of the bias from a change in the composition of enrollees induced by the change in cost sharing. Several recent US studies (Chandra et al., 2010a; Chandra et al., 2010b; Chandra et al., 2014) have used a quasi-experimental design by exploiting a change in the copayments of one health insurance plan and using unchanged insurance plans as a control group. However, the change in cost sharing could also affect people's decision to enroll in insurance plans. Such self-selection behavior could bias the elasticity estimates. For example, a larger proportion of people with less price sensitivity may continue their enrollment after the a cost-sharing increase, which may bias the price elasticity estimates toward zero. The Taiwanese National Health Insurance (NHI) is a single-payer scheme and every citizen is required to join the program.[71] Thus, our elasticity estimates are free of bias from change in the composition of the enrollees after the cost-sharing change.

[71] The only exceptions are citizens who lose their citizenship, die or are missing for more than six months.

Second, our paper provides credible and transparent estimates of price elasticities of healthcare utilization for young children, which is largely unexplored in the previous literature. Our RD design offers a unique opportunity to obtain estimates in a local randomized experiment. The comparison at the 3rd birthday convincingly isolates the impact of patient cost sharing on healthcare utilization from other factors because there are no confounding factors at the 3rd birthday.[72]

In addition, the data we use in this paper is administrative insurance claim data that contains all NHI records of healthcare payments and use for children under 4 years of age in Taiwan during our sample period.[73] Compared with survey data, administrative data have a number of advantages, such as much less measurement error and larger sample sizes. These features also allow us to get an accurate measure of the key variable in this paper — patient's age at visit. Prior studies using survey data find that there is substantial heaping in the reported birth dates of patients, which might reflect measurement error of patients' ages.[74] Also, our large sample size allows us to get precise estimates of the heterogeneity in the cost-sharing effect across different subgroups or types of healthcare that could not be analyzed precisely in the RAND HIE because of its limited sample of children.

Finally, this paper presents the first evidence on the effect of differential copayments for outpatient care on the choice of healthcare providers in Taiwan. NHI sets higher copayments for the visits to hospitals than those to clinics. This is because there is no gatekeeper system in Taiwan. Patients can choose healthcare providers freely without referral from primary care physicians. This freedom of choice might make some patients whose illness can be treated in clinics overuse outpatient resource in hospital (i.e. moral hazard) and crowd out those who have to get treatments in hospitals.

[72] In Taiwan, turning age 3 does not coincide with any confounding factors, such as age of starting school or a recommended immunization schedule. We will discuss this issue in Section 3.4.

[73] 99% of the Taiwanese population is covered by NHI. Furthermore, NHI covers almost all medical services. We will discuss this issue in more detail later.

[74] Shigeoka (2014) finds that respondents in Japanese Patient Survey tend to report the first day of the month as their birthday when they forget the exact date of birth.

NHI uses differential copayment to lead patients to choose their healthcare providers based on severity of illness and allocate outpatient resource in hospitals to the patients who need them most. We examine the effectiveness of differential copayments by comparing patient's choice of providers before the 3rd birthday (i.e. no copayment) and after the 3rd birthday (i.e. differential copayment).

The rest of the paper is organized as follows. Section 3.2 gives a brief overview of the institutional background. In Section 3.3, we discuss our data and sample selection. Section 3.4 describes our empirical strategy. In Section 3.5, we analyze the main results. Section 3.6 provides concluding remarks.

3.2 Policy Background

3.2.1 National Health Insurance in Taiwan

In March 1995, Taiwan established the NHI, which is a government-run, single-payer scheme administered by the Bureau of National Health Insurance. Prior to this, health insurance was provided through three main occupational forms — labor insurance for private-sector workers, government employee insurance, and farmers' insurance. These systems accounted for only 57% of the Taiwanese population (Lien et al., 2008). The remainder of the population were not employed: people over 65, children under 14, and unemployed workers. The implementation of the NHI raised the coverage rate of health insurance sharply to 92% by the end of 1995, and since 2000, it has stayed above 99%.

The NHI provides universal insurance coverage, with almost all medical services covered, such as outpatient, inpatient, dental, and mental health services, prescription drugs, and even traditional Chinese medicine. The NHI classifies healthcare providers into four categories based on accreditation: major teaching hospitals, minor teaching hospitals, community hospitals, and clinics.[75] As in most Asian countries, enrollees are free to choose

[75] The clinic is similar to the physician's office in Canada and the US.

their care providers and do not need to go through a general practitioner (i.e. family physician) to obtain a referral. For example, patients can directly access specialists in a major teaching hospital without a referral. In other words, the NHI does not adopt a gatekeeper system.[76]

3.2.2 Patient Cost-Sharing

Patient cost sharing in Taiwan comprises two parts: (1) the copayment (coinsurance);[77] (2) other non-NHI-covered medical costs (e.g. a registration fee for an outpatient visit).[78]

Cost-Sharing for Outpatient Care

With respect to outpatient care, a patient pays a copayment plus a registration fee for each visit.[79] If a physician prescribes a drug at a visit and the drug cost is above 100 NTD, the patient also needs to pay a share of the cost of the prescription drug, which is 20% of total drug cost. However, most visits for children under age 3 have drug costs below 100 NTD so patients usually do not pay for their prescription drug.[80] Compared with the copayment, the average out-of-pocket cost for outpatient prescription drugs (under age 3) is quite small, at only 2.5 NTD per visit.[81]

The copayments are based on a national fee schedule. In general, a higher copayment is set for the health providers that have higher accreditation.[82]

[76] For example, the National Health Service (NHS) in the United Kingdom adopts a gatekeeper system. Patients cannot directly obtain outpatient services at hospitals. Instead, they need to get a referral from a general practitioner. Provincial Health Insurance in Canada adopts a similar system.

[77] A copayment is a fixed fee paid by the insurance enrollee each time a medical service is accessed. Coinsurance is a percentage of the medical payment that the insured person has to pay. The NHI adopts copayments for outpatient care and coinsurance for outpatient prescription drugs and inpatient care.

[78] More discretionary healthcare, such as plastic surgery, sex reassignment surgery and assisted reproductive technology, etc., are not covered by the NHI. Patients have to pay the full cost for these services.

[79] Both are fixed amounts.

[80] If drug cost is under 100 NTD, a patient has no out-of-pocket cost.

[81] The average drug cost per visit is only 61 NTD, which is under 100 NTD and thus, patients do not pay any out-of-pocket cost at most visits.

[82] The NHI in Korea has a similar cost-sharing policy. Patients have to pay 40–50% of

The first rows of Panel A in Table 3.1 summarize the copayments for four types of providers during our sample period (2005 to 2008). To treat the same illness, a major teaching hospital charge a patient a copayment of 360 NTD (i.e. 11 USD) per outpatient visit but the copayment for one clinic visit is only 50 NTD (1.5 USD).[83]

The spirit of this design is to use the differential copayments to guide patients to properly choose their health providers based on the severity of an illness so as to better allocate medical resources to the patients who need them most. This design is needed because patients in Taiwan (and other Asian countries) have no restrictions on their choice of healthcare providers. For identical diagnosis, patients might get better and more treatments from the outpatient care in hospital. For example, patients could choose a teaching hospital or a clinic to treat a cold. But a teaching hospital could provide more tests and prescribe better drugs than a clinic. Table 3.5 compares major teaching hospitals, minor teaching hospitals, community hospitals, and clinics in terms of composition of medical expenses and composition of patients. In general, a patient pays higher share of medical expense for one visit to a teaching hospital than a clinic. A teaching hospital also provides more medical service than a clinic. The average medical expenses per teaching hospital visit is 987 NTD, which is three times as much as that per clinic visit. This is because a physician in a teaching hospital can conduct more health examinations (e.g. X-ray inspection) and medical treatments (e.g. therapeutic radiology) than one in a clinic. If there was no difference in the level of patient cost sharing between teaching hospitals and clinics, patients might overuse the limited medical resources of the hospitals and crowd out other patients whose illnesses could only be treated at hospitals.

In addition to the copayment, the patient must also pay a registration fee for each outpatient visit, which is not covered by the NHI. The registration fee reflects the health provider's administrative costs and is determined by

total medical costs when visiting hospitals but only 15–30% when visiting clinics.

[83]For more detailed information about the copayment schedule, please see the note in Table 3.1. A reimbursement is also paid according to the provider's accreditation. That is, major teaching hospitals can obtain the highest reimbursement for their medical services.

the provider.[84]

Cost-Sharing for Inpatient Care

For inpatient admissions, the patient cost sharing takes place through coinsurance. Depending on the length of the stay and the type of admission (acute or chronic admission), the coinsurance rate is 10% to 30% of the total medical costs per admission. For example, a patient must pay 10% of the hospitalization costs for the first 30 days they stay in an acute admission unit and 20% for the next 30 days. Almost all inpatient admissions for young children (99.5%) are acute admissions and the length of a stay in our sample is always within 30 days.[85] Thus, coinsurance rates for most admissions are around 10%. Panel B in Table 3.1 lists the coinsurance rates for inpatient services.[86]

Because inpatient care usually results in larger financial risks than outpatient care, the NHI has a stop-loss policy (i.e. maximum out-of-pocket cost) for inpatient admissions. The out-of-pocket cost must be no greater than the stop-loss, which is calculated annually as 10% of the gross domestic product per capita in Taiwan. The NHI covers all costs above the stop-loss.[87] According to NHI statistics, very few patients (less than 1%) reach this stop-loss, so the non-linearity imposed by it should not seriously bias our estimates of price elasticity.[88] Moreover, in contrast to health insurance plans in the US and other countries, the NHI does not require patients to

[84]Our main dataset lacks this information. However, the NHI has another database that provides information about the registration fees of all health providers during our sample period (2005–2008). Major teaching hospitals usually charge 150 NTD, minor and community hospitals 100 NTD, and clinics 50 NTD. We use this information to impute the registration fees for the four types of providers.

[85]In our empirical analysis, we limit our estimated sample for inpatient services to the cases with acute admissions with length of stay within 30 days.

[86]Some parents might buy private health insurance for their children. Such insurance can cover the out-of-pocket costs of inpatient care. Nevertheless, private health insurance for young children is not popular in Taiwan.

[87]In 2008, the annual maximum out-of-pocket cost is about 50,000 NTD.

[88]This is because the NHI waives the cost-sharing for patients with catastrophic illnesses (e.g. cancer), who would have a greater probability of reaching the stop-loss if their cost sharing were not waived.

pay deductibles before insurance coverage begins. The above two features substantially simplify our computation of the price elasticities.[89]

3.2.3 Change in Patient Cost Sharing at the 3rd Birthday

To reduce the financial burden on parents and ensure that every child obtains essential medical treatment in her early childhood, in March 2002, the Taiwan government enacted the Taiwan Children's Medical Subsidy Program (TCMSP). This program, through subsidies, exempts all copayments and coinsurance for outpatient visits, outpatient prescription drugs, inpatient admissions, and emergency room visits for children under the age of 3. A patient loses eligibility for subsidies at her 3rd birthday. Since the implementation of TCMSP, a patient under 3 years of age has only had to pay the medical costs not covered by the NHI (e.g. the registration fee for outpatient care and other non-covered medical services).[90]

Figure 3.1 plots the observed age profile of average out-of-pocket cost per outpatient visit and that of average out-of-pocket cost per inpatient admission (180 days before and after the 3rd birthday).[91] Figures A.3 and 3.1b reveal that patients experience a sharp increase in price for both outpatient and inpatient services at their 3rd birthday. Especially for inpatient services, the out-of-pocket cost per admission suddenly rises from zero to almost 1,300 NTD, which could bring about sizeable financial risk to a household with young children turning 3 years old.[92]

Note that the observed price changes per visit at the 3rd birthday are

[89] In health insurance, the deductible is the amount that an insured person has to pay before an insurer (e.g. the insurance company) starts to pay.

[90] If they use medical services not covered by the NHI, they will have to pay all expenses. However, the NHI does cover most health services. Those that are not covered are mostly quite discretionary, such as plastic surgery, sex reassignment surgery and assisted reproductive technology, etc.

[91] Each dot represents the ten days average price of each outpatient visit (inpatient admission) at a given age. The line is obtained by fitting a linear regression to the age variables fully interacted with a dummy indicating whether the child is age 3 or older.

[92] The average wage rate is around 225 NTD (i.e. 7 USD) per hour in 2006. The average monthly household earned income is around 45,000 NTD (or roughly 1,400 USD) in 2006. Therefore, out-of-pocket costs for an inpatient admission can account for 3 percent of a household's average monthly income.

endogenous. Especially for outpatient services, the price change at the 3rd birthday is larger for visits to a teaching hospital than to a clinic or community hospital. For example, the price per visit for a major teaching hospital increases by 240% (from 150 to 510 NTD) at the 3rd birthday and the price for a minor teaching hospital rises by 240% (from 100 to 340 NTD). However, the visit price for a clinic only increases by 100% (from 50 to 100 NTD). In other words, the TCMSP indeed subsidizes outpatient services in teaching hospitals much more than those in clinics or community hospitals. Therefore, patients might also change their choices of providers at their 3rd birthday, which could make the observed out-of-pocket cost per visit after the 3rd birthday endogenous (i.e. already reflected in the change in choice of provider). To obtain the exogenous price change at the 3rd birthday, we need to fix the utilization of each type of provider.

Table 3.2 presents the weighted average out-of-pocket cost per visit before and after the 3rd birthday.[93] The weights are the average daily utilization of each type of provider 90 days before the 3rd birthday. Thus, the numbers in the first row are the actual weighted average out-of-pocket costs per visit *before* the 3rd birthday and the numbers in the second row are counterfactual weighted average out-of-pocket costs per visit *after* the 3rd birthday, which uses the share of utilization of providers at age 2 (i.e. 90 days before the 3rd birthday) as weights. In this way, we can compute the difference between rows (1) and (2) to obtain the exogenous change in out-of-pocket costs per visit/admission at the 3rd birthday. Table 3.2 shows that the average price of outpatient visits rises by more than 100% (from 58.9 to 132.7 NTD) at the 3rd birthday, and the average price of inpatient admissions jumps sharply from zero to 1296 NTD. To sum up, in terms of both the level and the percentage change, the out-of-pocket cost for each inpatient admission sees a much larger increase than that for each outpatient visit.

[93] The bandwidth is 90 days. Thus, we use out-of-pocket cost per visit/admission within the 90 days before and after the 3rd birthday to obtain the estimates in Table 3.2.

3.3 Data and Sample

3.3.1 Data

To implement our empirical analysis, we need the following information: (1) the enrollee's exact age to the day at the time of a visit;[94] (2) the utilization of the outpatient or inpatient services; (3) the expenditures of the outpatient or inpatient services. We use unique claims data from Taiwan's National Health Insurance Research Database (NHIRD), which contains detailed information about patient's out-of-pocket costs, total healthcare expenditures and healthcare utilization for each outpatient visit (inpatient admission) of all NHI enrollees in Taiwan.[95] In addition, the NHIRD includes the exact dates of outpatient visits (inpatient admissions) and the exact birth date of every enrollee, which allows us to precisely measure the children's ages in days for our RD design.

For our purposes, we linked information from four types of files in the NHIRD: outpatient claims files, inpatient claims files, enrollment files, and provider files. First, outpatient (inpatient) claims files record information about payments and medical treatments for each visit. These files contain the enrollee's ID and birth date, the hospital or clinic ID, the date of the visit, the total healthcare expenditures, total out-of-pocket costs, diagnosis[96], and medical treatment.[97] Second, we use the enrollee's ID to merge the enrollment files and obtain each enrollee's demographic information, such as gender, household's monthly income, number of siblings, and town of residence. Finally, we use the hospital or clinic ID to link with the information (e.g. provider's accreditation) in the provider files.

[94] That is, we measure age in days.

[95] Due to privacy concern, NHIRD only allows at most 10% sampling for each research application. Thus, we only use claims data of sample with age 2 and 3 during 2005–2008 and 1997–2001.

[96] Diagnoses are recorded in five digits according to the ICD9 (International Classification of Diseases, Ninth Revision, Clinical Modification).

[97] Inpatient claims files have further information about length of stay.

3.3.2 Sample

To avoid the effect of variation in the cohort size on our estimation, we focus on the healthcare use from the same cohort (fixed panel). Our original sample is all NHI enrollees born between 2003 and 2004. The original sample size is 435,206 (see Table 3.3).[98] We further restrict our sample to those enrollees who were continuously registered in the NHI while aged 2 and 3, which reduces the sample size by 8,619. In addition, we eliminate those enrollees in the sample with cost-sharing waivers, such as children with catastrophic illnesses and children from very low-income families, since these children would not experience any price change when turning 3. The above procedure reduces our original sample by 5.7%, making the final sample size for estimation 410,517. Table 3.3 provides summary statistics of the characteristics of the enrollees at age 3, in the original sample and the final sample used in our empirical analysis. We find that the selected characteristics are quite similar between the two samples.

We use 2005–2008 NHIRD data to obtain all records of outpatient visits and inpatient admissions of these children when aged 2 or 3.[99] Following Lien et al. (2008), we also exclude visits relating to dental services, Chinese medicine, and health check-ups with a copayment waiver.[100]

Table 3.4 provides the descriptive statistics for the outpatient visits and inpatient admissions and compares their characteristics within 90 days before and after the 3rd birthday.[101] We find that children use more outpatient and inpatient care before their 3rd birthday. Most young children visit clinics for outpatient services. However, they tend to visit teaching hospitals more frequently before their 3rd birthday than after it.

[98]Since 99% of Taiwanese are covered by the NHI, this sample represents nearly the entire population of children born between 2003 and 2004 in Taiwan.

[99]The sample period was chosen because children born in 2003 are aged 2 in 2005–2006 and children born in 2004 are aged 3 in 2007–2008.

[100]The NHI provides nine health check-ups with copayment waiver for children under the age of 7. Since patient cost sharing for these visits does not change at the 3rd birthday, we eliminate them to avoid biased estimations.

[101]We make this choice because our main results use 90 days as the bandwidth.

3.4 Empirical Specification

Our identification strategy is similar to that in recent studies utilizing an "age discontinuity" to identify the insurance coverage effect (Card et al., 2008; Card et al., 2009; Anderson et al., 2012) or patient cost-sharing effect (Shigeoka, 2014) on medical utilization by adults or the elderly. We are the first to apply the RD design to study the impact of patient cost sharing on healthcare utilization and expenditure for young children. The general form of our RD regression is as follows:

$$Y_i = \beta_0 + \beta_1 Age3_i + f(a_i;\gamma) + \varepsilon_i \qquad (3.1)$$

where Y_i is the outcome of interest for the child i, namely (1) the number of outpatient visits or inpatient admissions; (2) the total expenditure of outpatient or inpatient care; (3) the expenditure per outpatient visit (inpatient admission) at a given age. The variable a_i is child i's age and is measured in days. The variable $Age3_i$ is a treatment dummy that captures the higher level of patient cost sharing (i.e. loss of cost-sharing subsidy) at the 3rd birthday and is equal to one if child i is age 3 or older. The 3rd birthday is the 1096th or 1095th day after birth.[102] The key assumption of the RD design is that the age profile of the healthcare utilization is smooth (continuous). Thus, we assume $f(a_i;\gamma)$ to be a smooth function of age with parameter vector γ that accommodates the age profile of the outcome variables. The ε_i is an error term that reflects all the other factors that affect the outcome variables. Our primary interest is β_1, that measures any deviation from the continuous relation between age and the outcomes Y_i at child i's 3rd birthday (i.e. when the treatment variable switches from 0 to 1). If no other factors change discontinuously around the child's 3rd birthday, that is, $E[\varepsilon_i|a_i]$ is continuous at age 3, β_1 represents the causal effect of the higher level of patient cost sharing on the expenditure and on utilization of

[102] Since 2004 is leap year, its February has 29 days. For the children born before 2004 February 29th, their 3rd birthday would be 1096th day after birth (365 x 3 + 1 = 1096). For those born after 2004 March 1st, their 3rd birthday would be 1095th day after birth.

young children's healthcare. In general, there are two ways to estimate β_1, typically referred to as the global polynomial approach and the local linear approach (Lee and Lemieux, 2010).

In the global polynomial approach, we can use all available data to capture the age profile of healthcare utilization $f(a_i;\gamma)$ by using a flexible parametric function (e.g. in our analysis we use a third-order polynomial of age).[103] One caveat of this approach is that an incorrect functional form for the regression could create a biased estimate of β_1. To avoid a misspecification bias, we adopt a local linear regression as our main specification and present the global polynomial estimates for comparison.

In the local linear approach, we capture the age trend of the healthcare use $f(a_i;\gamma)$ by estimating a linear function over a specific narrow range of data on either side of the threshold (i.e. 3rd birthday). The local linear estimates of the treatment effect are the differences between the estimated limits of the outcome variables on each side of the discontinuity. Our baseline specification is the following local linear regression:

$$Y_i = \beta_0 + \beta_1 Age3_i + \gamma_1(a_i - 1096) + \gamma_2 Age3_i(a_i - 1096) + \varepsilon_i \qquad (3.2)$$

In practice, we obtain the estimated treatment effect β_1 by allowing the slope of the age profile to be different on either side of the 3rd birthday, by interacting the age variable fully with the intercept and $Age3_i$. Also, we recenter the age variable to the 3rd birthday to make β_1 directly represent the treatment effect at the 3rd birthday.[104] The equation (3.2) is estimated via weighted least squares using a triangular kernel (i.e. giving more weight to the data points close to the 3rd birthday). We restrict our sample to the 90 days before and after the 3rd birthday. The choice of bandwidth and the computation of the standard errors of the discontinuity estimates are important issues for local linear estimation. In Table B.3, we show that

[103] We have all NHI records of medical utilization within 365 days before and after each individual's 3rd birthday (i.e. from 2nd birthday to their 4th birthday).

[104] For the children born before 2004 February 29th, age variable is $a_i - 1096$. For those born after 2004 March 1st, age variable is $a_i - 1095$.

our main estimates are robust to various choices of bandwidth and different methods of calculating the standard errors.[105]

Following Card et al. (2009), Anderson et al. (2012) and Lemieux and Milligan (2008), we collapse the individual-level data into age cells (measured in days), which gives us the same estimates as the results from the individual-level data but substantially reduces the computational burden. Therefore, our regressions are estimated on day-level means for each day of age:

$$Y_a = \beta_0 + \beta_1 Age3 + \gamma_1(a - 1096) + \gamma_2 Age3(a - 1096) + \varepsilon_a \qquad (3.3)$$

We also take logs of our dependent variables to allow β_1 to be interpreted as the percentage change in the dependent variables. That is, the dependent variables for the RD estimation are the log of total outpatient (inpatient) expenditure, the log of the total number of outpatient visits (inpatient admissions), and the log of outpatient (inpatient) expenditure per visit, at each day of age. The most important assumption for our RD estimation is that, except for the higher level of patient cost sharing, there is no change in any other confounding factors that affect the demand for healthcare services at

[105] Deciding how "narrow" a range of data to use, namely, choice of bandwidth, is critical to local linear estimation. If the bandwidth were too wide, the local linear estimate β_1 could be biased due to misspecification. That is, the linear function would be unable to capture the age profile over such a "wide" range of data. If the bandwidth were too narrow, there would not be enough data for the estimation to get a precise local linear estimate. Thus, the optimal bandwidth needs to balance bias and precision (variance) to estimate β_1. This is quite an active field in the nonparametric literature and there are many competing methods of selecting the optimal bandwidth, such as the plug-in approach (Imbens and Kalyanaraman, 2012; Cattaneo et al., 2013) and the cross-validation approach (Ludwig and Miller, 2007). In Table B.3, we show that our main estimates are robust across various optimal bandwidth selectors. In addition, the standard error of the discontinuity estimate is an important issue in local linear estimation since the available bandwidth selectors tend to give a "large" bandwidth and lead to biased local linear estimates. One solution is to use bias-correction estimates. However, the conventional standard error of the bias-correction estimates fails to consider the variability of additional second-order bias estimates, which results in standard errors that are too small and false statistical inferences. Cattaneo et al. (2013) proposes a method of accounting for this variability to obtain the robust standard error and confidence interval. In Table B.3, we show that the statistical inferences of our main estimates are still valid even if we use more conservative way to compute our standard error.

the 3rd birthday. For this age group, potential confounding factors could include vaccination and pre-school attendance. The recommended immunization schedule could mechanically increase the healthcare spending and use of young children at age 3. However, this concern is alleviated since children in Taiwan do not need to have vaccines at age 3 and indeed take most vaccines before they are 2 years of age (Center of Disease and Control, 2013).[106] On the other hand, entering pre-school could increase the chance of a child picking up illnesses (e.g. the flu), which would affect children's healthcare use. This factor might not interfere with the cost-sharing change at age 3 because the age of entry for "public" pre-schools is 4 years of age and the government does not specify a statutory attendance age for "private" kindergartens. Most importantly, we measure the children's age at a daily level, so our RD design will be invalid only if these factors also change abruptly within one or two days of the 3rd birthday. This fact substantially alleviates the concern that our estimates could be biased by other factors. We conduct several placebo tests to further confirm the validity of our RD design (e.g. using data before 2002 when TCMSP was implemented).

3.5 Results

In this section, we examine the impact of the higher cost sharing at a child's 3rd birthday on healthcare expenditure and utilization. As mentioned above, our sample consists of the children born between 2003 and 2004 who were continuously enrolled in the NHI over the ages of 2 and 3. We follow these individuals across their 3rd birthdays to estimate the change in healthcare utilization and expenditure at age 3. We will examine outpatient care first and then impatient care.

3.5.1 Outpatient Visits and Expenditure

From Section 3.2, we know that the average out-of-pocket cost for each outpatient visit increases by more than 100% when a child passes their

[106] http://www.cdc.gov.tw/professional/page.aspx?treeid=5B0231BEB94EDFFC&nowtreeid=1B4BACA0D1FDDB84

3rd birthday. Our main question is how children's healthcare utilization and expenditure respond to this exogenous price change. We begin with a graphical analysis.

Graphical Analysis

Figure 3.2a shows the actual and fitted age profiles of total outpatient expenditure for children born between 2003 and 2004. The dots in the figure represent total outpatient expenditure per 10,000 person-years by patient's age at each visit (measured in days).[107] The solid line shows the fitted values from a local linear regression that interacts intercept and the age variables fully with a dummy indicating that the child has passed her 3rd birthday.[108] Corresponding to a sharp increase in patient cost sharing at the 3rd birthday, there is an obvious discrete reduction in outpatient expenditure when the children turn 3. The change in total outpatient expenditure can be decomposed into the change in the number of visits and the outpatient expenditure per visit. Figures 3.2c and 3.2e represent the actual and fitted age profiles of outpatient visits per 10,000 person-years[109] and outpatient expenditure per visit, respectively. We find that both variables also suddenly jump down, right after the children's 3rd birthday. On the other hand, we use pre-reform data (1997–2001) to plot the related outcome variables in Figures 3.2b, 3.2d and 3.2f. In sharp contrast to the graphs presented above, We find no visible discontinuity at the 3rd birthday.

Main Results

Table 3.6 presents the estimated impact of the 3rd birthday on outpatient expenditure and visits before (1997–2001) and after (2005–2008) the TCMSP

[107] We compute the total outpatient expenditure per 10,000 person-years by dividing the total outpatient expenditure at a particular age by the number of enrollees born between 2003 and 2004 and then multiplying this by 10,000. This is a common way to present data in the health economics and public health literatures and helps us to compare the estimated results across different sample periods and subgroups. Each dot represents 10-days average of the dependent variable.

[108] We use 90 days as our bandwidth.

[109] Again, each dot represents outpatient visits per 10,000 person-years at a given age, averaged over 10 days.

was introduced. Each panel displays results for different dependent variables of interest. Odd-numbered columns present RD estimates from a nonparametric local linear regression and even-numbered columns present RD estimates from a parametric OLS regression (cubic spline). Column (1) of Table 3.6 presents our main results for outpatient services and displays the estimates from a local linear regression with a triangular kernel function and a bandwidth of 90 days of age.[110] Corresponding to the sharp drop in outpatient expenditure at the 3rd birthday in Figure 3.2a, Panel A shows that the rise in the level of patient cost-sharing at the 3rd birthday causes overall outpatient expenditure to decrease significantly by 6.9%. The implied price elasticity of outpatient expenditure is around -0.10.[111]

The change in total outpatient expenditure comes from two margins: (1) the number of visits (extensive margin); (2) the outpatient expenditure per visit (intensive margin). Panel B reveals that the number of outpatient visits decreases by 4.7% at the 3rd birthday, which is smaller than the change in total expenditure. The remaining change comes from the change in the outpatient expenditures per visit. Panel C reveals that the outpatient expenditure per visit decreases significantly, by 2.2%, at the 3rd birthday. In fact, this result is likely to be a combination of two forces. First, higher cost sharing at the 3rd birthday could change the composition of patients and result in higher outpatient expenditure per visit at age 3. Assuming that the marginal patients are not as sick as those who use healthcare service regardless of cost-sharing subsidy eligibility, the average health of the patients may drop discretely at the 3rd birthday, leading to higher expenditures per

[110] We only use observation whose age at each visit is within 90 days before and after the 3rd birthday.

[111] This elasticity is calculated in the form of price elasticity. The standard formula for the price elasticity of demand is $((Q_2 - Q_1)/Q_1)/((P_2 - P_1)/P_1)$, where Q_1 and P_1 denote the baseline healthcare demand and patient cost sharing, respectively, and Q_2 and P_2 are the healthcare demand and patient cost sharing after the change in cost sharing. However, in the health economics literature, many studies (Leibowitz et al., 1985; Manning et al., 1981; Chandra et al., 2010a) also use the price elasticity, which denotes the percentage change relative to the average, since P_1 could be zero in some cases (e.g. the free plan in Rand HIE or zero out-of-pocket cost for inpatient care in this paper) and then the denominator of the price elasticity would be undefined. That is, the price elasticity is calculated as $((Q_2 - Q_1)/((Q_1 + Q_2)/2))/((P_2 - P_1)/((P_1 + P_2)/2)$.

visit.[112] Second, losing the cost-sharing subsidy at the 3rd birthday could also affect a patient's choice of provider (quality of each visit) and lead to lower outpatient expenditure per visit at age 3. As mentioned in Section 3.2, TCMSP indeed subsidizes more out-of-pocket costs for teaching hospital patients than clinic and community hospital patients, which would encourage patients to use outpatient services at teaching hospitals before the 3rd birthday, as patients could thereby extract greater subsidies but also receive a better quality of medical service.[113] Therefore, when patients lose their eligibility for the cost-sharing subsidy at the 3rd birthday, they may reduce their visits to teaching hospitals, resulting in lower expenditures per visit.[114] Our estimates in Panel C imply that the latter force dominates the former, causing outpatient expenditure per visit to exhibit a discrete drop at the 3rd birthday. In a later section, we will discuss this issue in more detail.

Validity and Robustness Checks

Columns (3) and (4) in Table 3.6 display the results of a placebo test using pre-reform data (1997–2001). The results reveal that there is no discontinuity in our outcome variables at the 3rd birthday before 2002 (when TCMSP was introduced). The point estimates are insignificant and close to zero, which substantially reduces concerns about the impact of other confounding factors on our estimates. In Table B.1, we conduct another placebo test by examining any discontinuities at other age cut-offs. We find our outcome variables (log of outpatient expenditure and number of visits) to be smooth across all selected age cut-offs, except for the 3rd birthday (i.e. 1096 days old).[115]

For a robustness check of our main specification, we use an alternative

[112]This assumes that healthcare providers spend more on treating less healthy patients.
[113]Every three to four years, the Ministry of Health and Welfare evaluates every NHI-contracted hospital/clinic to determine their accreditation. The category of "major teaching hospital" is seen as indicating the best-quality providers.
[114]Because the teaching hospitals may provide more medical services at each visit, such as health checks or medical treatments, it will cost more for each visit.
[115]There are several "significant" discontinuities at other age cut-offs. However, their magnitudes are quite small.

method (global polynomial approach) to estimate the discontinuity in the outcome variables at the 3rd birthday using all available data (365 days before and after the 3rd birthday) and a third-order polynomial age function with different slopes on either side of the 3rd birthday. Column (2) in Table 3.6 presents very similar estimates to our main results. In Table B. 2, we systematically examine the sensitivity of our RD estimates to different bandwidths and orders of polynomial. The estimates are fairly stable across different specifications. In Table B.3, we present various local linear estimates from three different bandwidth selectors and kernel functions to show that our main results are robust to these choices.

One caveat could threaten the validity of our RD design. Because every child eventually "ages out" of her cost-sharing subsidy, parents may anticipate the sharp increase in the price of medical services after the child's 3rd birthday and "stock up" on outpatient care.[116] This behavioral response would represent an inter-temporal substitution of healthcare (i.e. substituting future healthcare with current healthcare) and not a "real" change (increase) in the demand for healthcare induced by the cost-sharing subsidy, which is our main interest. Such a behavioral response would tend to bias upward our estimates of the change in healthcare utilization at the 3rd birthday (i.e. the price elasticity of healthcare utilization). From Figures 3.2a and 3.2c, we indeed find that outpatient expenditure and visits suddenly rise 20 days before the 3rd birthday. In order to account for the possible anticipation effect, we conduct a "donut" RD (Barreca et al., 2011; Shigeoka, 2014) by systematically excluding outpatient expenditure and visits within 3–21 days before and after the 3rd birthday (Table B.4). Although there is no consensus on the optimal size of a donut hole, and while eliminating the sample around the threshold seems to contrast with the spirit of RD design, this type of estimation can still give us some sense of the "stocking up" effect on our estimates. Table B.4 indicates that the estimates from different sizes

[116] Since most outpatient visits of young children are for acute diseases (e.g. 74% of visits are for respiratory diseases), it is hard to believe that parents would be able to substitute children's outpatient care today for care in one month. However, it would be possible to substitute outpatient care within a few days.

of donut hole give us very similar results to our main RD estimates.

Change in Choice of Providers at 3rd birthday

The NHI in Taiwan (and other Asian countries) does not adopt a gatekeeper system to restrict patients' choices of providers. Instead, the NHI sets different levels of copayments for four different types of providers to encourage patients to choose the most suitable provider based on their understanding of the seriousness of the illness and to rectify possible moral hazard behaviors in choosing providers. As mentioned before, the TCMSP exempts all copayments for children under the age of 3, which gives us a unique opportunity to examine the impact of differential copayments on the patient's choice of provider by comparing the choices right before the 3rd birthday (i.e. no copayments) with those right after the 3rd birthday (i.e. differential copayments).[117]

Figures 3.3a to 3.3d present the age profiles of the outpatient visits by type of provider. We find that outpatient visits to major and minor teaching hospitals see strikingly discrete reductions just after the 3rd birthday. However, the number of visits to community hospitals exhibit the opposite pattern, namely jumping at the 3rd birthday, and there is a less obvious drop in visits to clinics after the 3rd birthday. Most of the decline in the overall number of outpatient visits indeed comes from the teaching hospitals. The visual evidence suggests that the change in relative prices at the 3rd birthday results in a significant redistribution of caseloads across different types of providers.

Coinciding with the graphical evidence, the RD estimates in Panel B of Table 3.7 show that turning age 3 substantially reduces the number of outpatient visits to major and minor teaching hospitals, by 59% and 44%, respectively. However, outpatient visits to community hospitals increase by 18% and the caseloads of clinics decrease only slightly, by 2%. This result indicates that patients are quite sensitive to the relative prices of different types of providers, and can switch providers easily. The question

[117] Before the 3rd birthday, patients still need to a pay registration fee. However, the registration fee does not vary substantially across different providers.

that follows is what kind of healthcare can easily be substituted between teaching hospitals and community hospitals/clinics ?

In Panel C of Table 3.7, we use outpatient expenditure per visit as a proxy for severity of illness.[118] The estimates in Panel C reveal that turning age 3 substantially increases the expenditure per visit to the major and minor teaching hospitals by 20% and 6%, respectively. This result implies that most of the reduction in visits to teaching hospitals at the 3rd birthday is actually related to less severe diseases.

Since patients switch their utilization from teaching hospitals to community hospitals/clinics right after the 3rd birthday, we suspect that the reduced visits to teaching hospitals could be for the illnesses for which it is not necessary to attend a teaching hospital. That is, they can be treated at clinics or community hospitals instead, which implies a substantial moral hazard whereby outpatient care in teaching hospitals are overused before the 3rd birthday when patients do not pay copayments. The above results suggest that the differing levels of copayments are an important factor in patients' choice of providers. Maintaining differential copayments between different types of providers could be an effective tool for allocating medical resources efficiently.[119]

Heterogeneous Effect

In this section, we investigate the heterogeneity of price responses across different types of diagnoses and various subgroups of young children. Each row displays a different type of diagnosis and subgroup. Column (1) in Table 3.8 presents the rate of outpatient visits per 10,000 person years 90 days before the 3rd birthday to give us some insights about the relative size of outpatient visits across different types of diagnoses and subgroups before a child's 3rd birthday. Column (2) and (3) in Table 3.8 display the RD

[118] Here we assume that more severe diseases would incur higher expenditures per visit.

[119] There are 24 major teaching hospitals and 65 minor teaching hospitals in Taiwan. Most of them are located in urban area. However, Taiwan is highly urbanized (e.g. around 78% people live in cities). In addition, Taiwan has a small geography so it would not take much time for a patient living in a rural area to reach the closest teaching hospital.

estimates of outpatient expenditure (taking logs) and implied price elasticity of expenditure, respectively. Panel A in Table 3.8 presents the results for selected diagnoses. The first three rows in Panel A list the top three common visit diagnoses for young children and all of them are acute respiratory diseases: upper respiratory infection (URI), acute bronchitis, and acute sinusitis, which accounts for 40% of total outpatient visits.[120] For some diseases, such as acute bronchitis and sinusitis, receiving proper outpatient care could be beneficial to children's health. Column (2) in Panel A shows that the outpatient expenditure for these common diagnoses significantly decline after the 3rd birthday. However, the estimated sizes of the reduction at the 3rd birthday for these diseases are smaller than estimates from overall outpatient expenditure. The implied price elasticities of expenditure are only -0.04 to -0.08, which reveals patients (parents) are not price sensitive to outpatient care for acute respiratory diseases.[121]

The remaining rows in Panel A present RD estimates for other selective diagnoses that may be less serious but need timely treatment to improve living quality, such as skin diseases. Losing the cost sharing subsidy causes a 14.9% reduction in outpatient expenditure for skin diseases, which is much larger than the overall decline in outpatient expenditure. Much larger decreases can also be found for outpatient care that are more discretionary but could reduce future healthcare costs, such as mental health service and preventive care. Turning three substantially reduces outpatient expenditure for mental illnesses by 23.2% and for preventive care by 24.5%. The implied price elasticities for this type of healthcare are quite large (-0.33 for mental health service and -0.69 for preventive care).[122] Our results suggest preventive and mental care are quite price sensitive, which is especially interesting since preventive care and early treatment for children's mental disorders (e.g. autism) could result in better treatment outcomes and might substantially reduce future medical costs. This result offers an evidence in support

[120] $(119 + 51 + 48)/542 = 0.40$

[121] We use the same method mentioned in section 3.2 to obtain an exogenous price change at the 3rd birthday for each disease and then calculate the price elasticity (price elasticity).

[122] We use the same method mentioned in section 3.2 to obtain an exogenous price change at 3rd birthday for each disease and then calculate the price elasticity.

of "behavioral hazard" suggesting patients reduce utilization of healthcare services that are potentially high-valued (Baicker et al., 2013).

Panels B to D in Table 3.8 examine the distinct price response across various subgroups of young children. Panel B displays the results by birth order. I find that the baseline visit rate for the two groups are quite similar and the price elasticity of healthcare expenditure for 1st born children is significantly smaller than that for non-1st born children (in absolute terms). This results implies patients may be more cautious when raising the first-born children. They might think healthcare services for 1st born children are more necessary. So they may be less willing to reduce 1st born children's healthcare utilization when facing higher outpatient price. Panel C presents the results by gender. We observe two facts. First, outpatient utilizations for male are more price sensitive than those for female. Second, before the 3rd birthday, males have more outpatient visits than females.

These results might reveal a son preference of parents. Parents take their boys to see a doctor more often than girls when these is a cost sharing subsidy. These visits could be discretionary and price sensitive. Therefore, we observe outpatient care utilization for sons has a larger response to a change in cost sharing at the 3rd birthday than that for daughters. Panel D presents the results based on household income.[123] This subgroup analysis can help us get some sense of the income effect on children's outpatient utilization. If an income effect plays an important role in patient's utilization decision, we should expect those unable to afford a healthcare services after the price increase, such as low-income children, to reduce their utilization of healthcare more at the 3rd birthday than high-income children. We find that there is no difference in price sensitivity between high and low-income children. This implies higher cost sharing does not result in larger reduction in outpatient utilization for low-income children, suggesting the income effect might play a limited role in explaining the decreased outpatient utilization at the 3rd

[123] A low-income household is defined as follows: A household with monthly household income below 40,000 NTD. A high-income refers a household with monthly household income above 40,001 NTD. I also use more elaborate income categories(i.e. five income categories) to allow effects to vary. The results are quite similar to Panel D in Table 3.8.

birthday.

3.5.2 Inpatient Admissions and Expenditures

For young children, inpatient admissions are much less common than outpatient visits. Among our sample at age 2, the average annual number of outpatient visits is 19.8 but the average annual number of inpatient admissions is only 0.14. Nevertheless, the cost to the patient of one inpatient admission is 29 times more than that per outpatient visit and 17% of healthcare spending for young children is attributed to inpatient care. More importantly, patient cost sharing for inpatient admissions experiences a much larger increase at the 3rd birthday than does that for outpatient visits, in terms of both the level and the percentage change.[124] That is, inpatient care could have substantial impacts on overall healthcare spending and individuals' out-of-pocket cost. Hence, understanding how young children's demand for inpatient care responds to cost sharing has important policy and welfare implications.

However, the effect of turning age 3 (losing the cost-sharing subsidy) on the utilization of inpatient care is theoretically ambiguous. On the one hand, children may have fewer inpatient admissions and lower expenditure after they turn 3 because the patient cost sharing for inpatient care increases sharply at the 3rd birthday. On the other hand, the type of inpatient care that young children usually use might not respond to the price change. Most admission diagnoses in early childhood, such as pneumonia and acute gastroenteritis, can be treated with medication or bed rest. Previous studies (Card et al., 2008; Shigeoka, 2014) have found that patient cost sharing (or insurance coverage) has less impact on this type of diagnosis for the elderly. In addition, for young children, admissions requiring surgery are seldom selective (e.g. osteoarthritis, hip and knee replacement) but more likely life threatening and essential (e.g. congenital heart disease). Thus, we should expect inpatient care for young children to be less sensitive to price changes

[124] Average patient cost sharing for one inpatient admission increases by 1296 NTD at the 3rd birthday. However, the average price for one outpatient visit rises by just 74 NTD.

at the 3rd birthday.

Graphical Analysis

Figure 3.6a shows the actual and fitted age profiles of inpatient admissions for children born between 2003 and 2004. Similar to the graphs for outpatient care (Figure 3.2), the markers represent total inpatient expenditure per 10,000 person-years at the given age, which is measured in days from the 3rd birthday. The solid line shows the predicted values from a local linear regression that interacts the age variables fully with intercept and a dummy indicating that the child has passed her or his 3rd birthday. Surprisingly, in contrast to the sharp drop in outpatient expenditure, Figure 3.6a shows that inpatient expenditure exhibits no change at the 3rd birthday. Similarly, Figures 3.6c and 3.6e represent the actual and predicted age profiles of inpatient admissions and inpatient expenditure per admission. We also find that there is little visual evidence of any discontinuity in either inpatient admissions or inpatient expenditure per admission at the 3rd birthday. When we compare these with the graphs plotted using pre-reform data (1997–2001), we find the outcome variables in the pre and post-reform periods to have very similar age profiles.

Main Results

Table 3.9 presents the estimated effect of the 3rd birthday on inpatient expenditure and admissions before (1997–2001) and after (2005–2008) the introduction of the TCMSP. As in Table 3.6 for outpatient services, each panel displays results for a different dependent variable of interest. Odd-numbered columns present the RD estimates from nonparametric local linear regressions and even-numbered columns present the RD estimates from parametric OLS regressions (cubic spline). Consistent with the graphical evidence in Figure 3.6, all RD specifications in Table 3.9 suggest there is no statistically significant impact of turning age 3 on inpatient expenditure and utilization. The point estimates in column (1) of Table 3.9 (our baseline estimation) are close to zero and insignificant. They reveal that losing the cost-sharing subsidy reduces the total inpatient expenditure by only 0.89% and the number of inpatient admissions by 0.18%. The implied price elasticity of inpatient expenditure is about -0.004.[125]

There is little evidence on the impact of patient cost sharing on the demand for inpatient services. Our results are consistent with the findings in the prior literature. Shigeoka (2014) finds that the demand for inpatient admissions treated with bed rest and medication do not respond to the price change at age 70 in Japan. Card et al. (2008) obtained similar findings for Medicare recipients in the US. Since most admissions for young children involve these types of inpatient care, our results suggest that the utilization of inpatient care for young children could have a very limited response to patient cost sharing, which implies that young children's demand for inpatient care may not be discretionary but necessary. In other words, full insurance coverage of children's inpatient care does not cause moral hazard but substantially reduces the financial risk brought about by inpatient admissions.

3.6 Conclusion

In this paper, we provide convincing evidence on the price response of healthcare for young children. We exploit a sharp increase in the required level of patient cost sharing at age 3 in Taiwan that occurs when young children "age out" of the cost-sharing subsidy, which results in a higher level of patient cost sharing for children just after their 3rd birthdays than just before. We apply an RD design to estimate the impact of cost sharing on healthcare utilization in early childhood. We reach three conclusions. First, the demand for outpatient services responds significantly to the change in copayments, but the estimated price elasticity of outpatient expenditure is modest (at around -0.10). Second, differential copayments for outpatient care between hospitals and clinics represent an effective policy tool for encouraging patients to visit suitable providers based on the seriousness of

[125] Again, it uses the price change in Table 3.2 and is calculated in the form of price elasticity.

their illness. According to our estimates, due to the differential copayments, the number of visits to teaching hospitals is reduced by 50% and most of the foregone visits are for less severe conditions, which can be also treated in clinics. Finally, the demand for inpatient services does not respond to the price change. The implied price elasticity of inpatient expenditure is close to zero. The Rand HIE found mixed evidence on this issue and could not draw strong conclusions. Our results largely support the view that inpatient care for young children is not price sensitive. Taken together, these results suggest that the level of patient cost sharing for young children should differ between healthcare services and healthcare providers. For example, our results imply providing full insurance coverage for children's inpatient care can substantially reduce the financial risk for the households but does not induce excessive utilization of inpatient care. On the other hand, our estimates suggest having a higher level of copayment for outpatient care at teaching hospitals can reduce patient's moral hazard behavior of choosing healthcare providers, namely, attending teaching hospitals when they do not need to do so.

Several important questions have not been analyzed in this paper, such as the long-run health impact of this cost-sharing subsidy program. Future research could focus on this issue and this would give us a more complete picture of the effect of similar programs around the world.

3.7 Figures

Figure 3.1: Age Profile of Out-of-Pocket Costs

(a) Average price per outpatient visit (NTD)

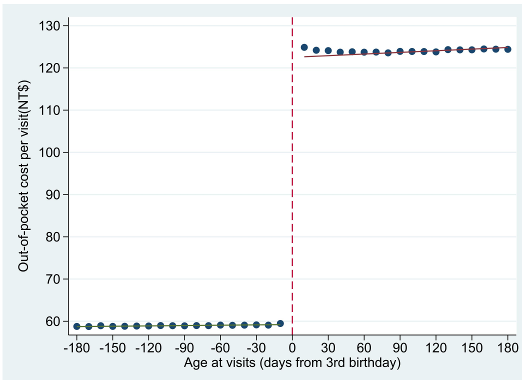

(b) Average price per inpatient admission (NTD)

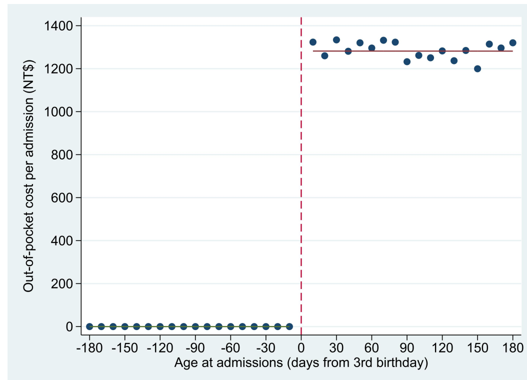

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday. The dependent variable is average price per outpatient visit (inpatient admission) by patient's age at visit (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

Figure 3.2: Age Profile of Outpatient Expenditure and Visits

(a) Outpatient expenditures per 10,000 person-years: 2005–2008

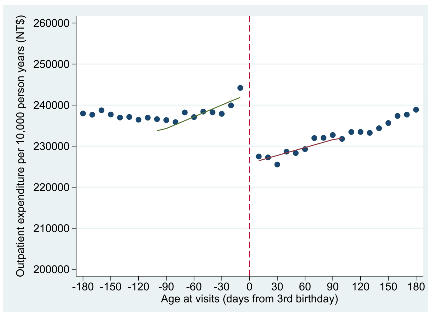

(b) Outpatient expenditures per 10,000 person-years: 1997–2001

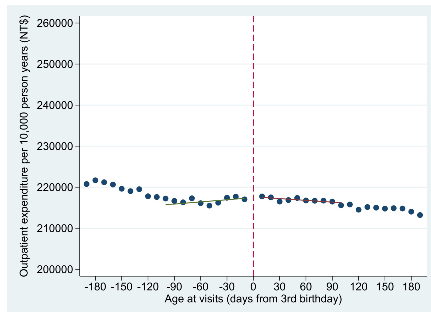

(c) Outpatient visits per 10,000 person-years: 2005–2008

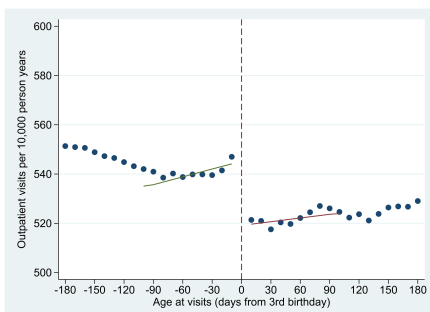

(d) Outpatient visits per 10,000 person-years: 1997–2001

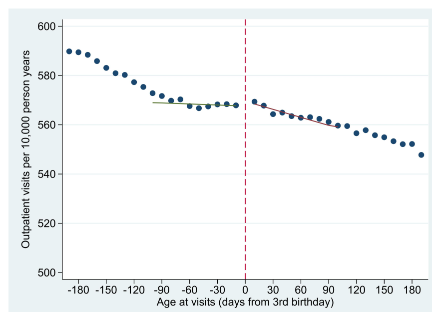

(e) Outpatient expenditures per visit : 2005–2008

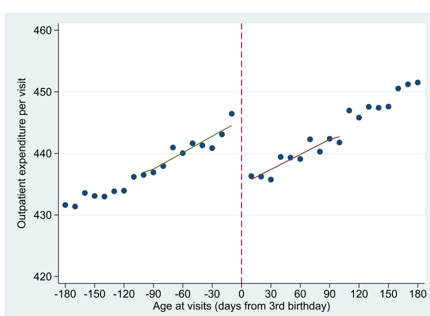

(f) Outpatient expenditures per visit: 1997–2001

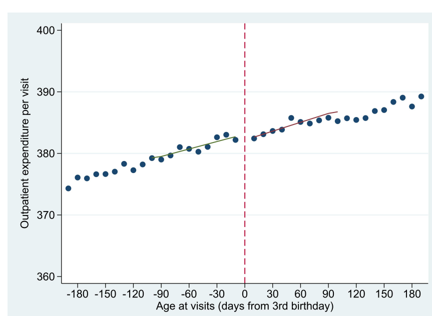

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday (90 days bandwidth). The dependent variables are outpatient expenditure per 10,000 person years, outpatient visits per 10,000 person years, and outpatient expenditure per visit by patient's age at visit (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

Figure 3.3: Age Profile of Outpatient Visits per 10,000 Person-Years by Type of Provider

(a) Major Teaching Hospital

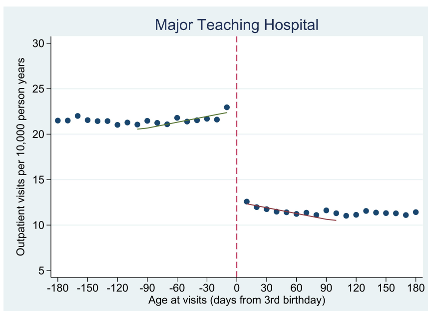

(b) Minor Teaching Hospital

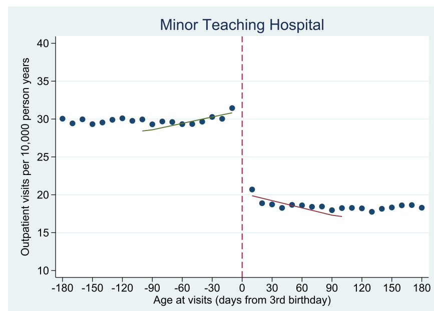

(c) Community Hospital

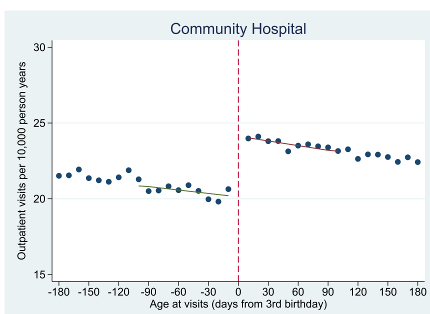

(d) Clinic

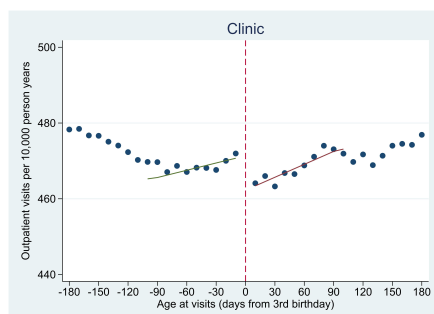

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday (90 days bandwidth). The dependent variables are outpatient visits per 10,000 person years (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

Figure 3.4: Age Profile of Outpatient Expenditure per 10,000 Person-Years (NTD) by Type of Provider

(a) Major Teaching Hospital

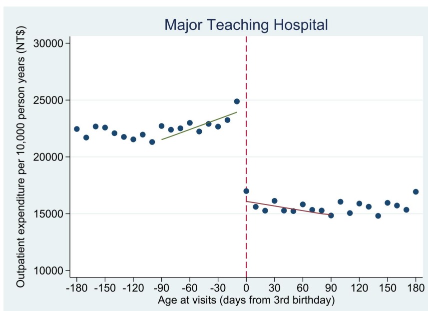

(b) Minor Teaching Hospital

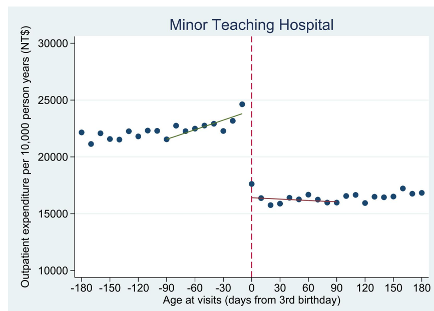

(c) Community Hospital

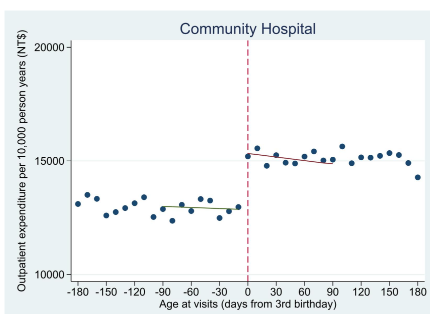

(d) Clinic

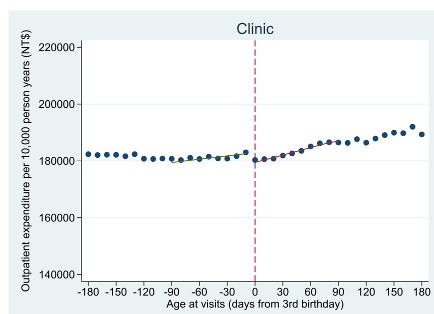

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday (90 days bandwidth). The dependent variables are outpatient visits per 10,000 person years (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

Figure 3.5: Age Profile of Outpatient Visits per 10,000 Person-Years by Diagnosis

(a) Acute upper respiratory infection

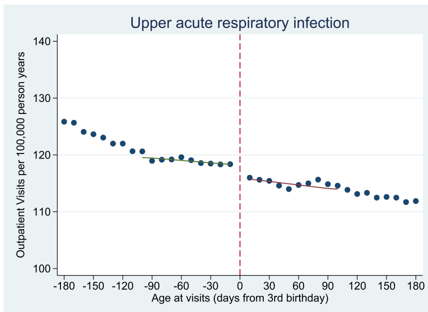

(b) Bronchitus

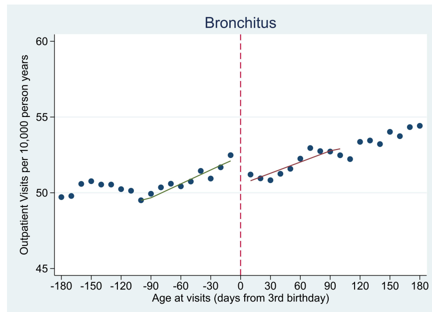

(c) Sinusitis

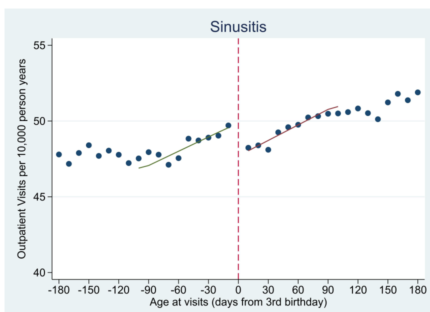

(d) Diseases of the skin

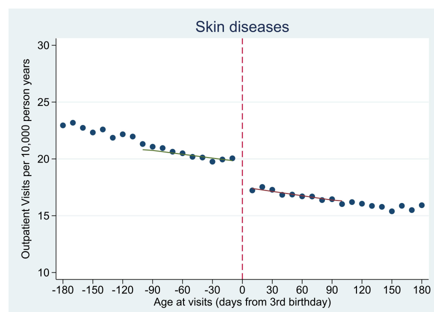

(e) Mental illnesses

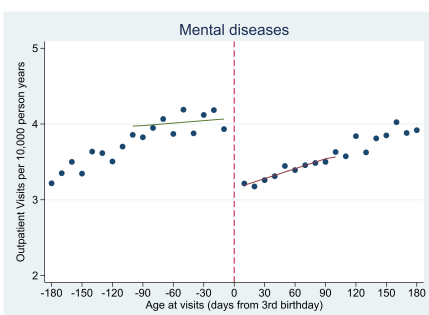

(f) Preventive care

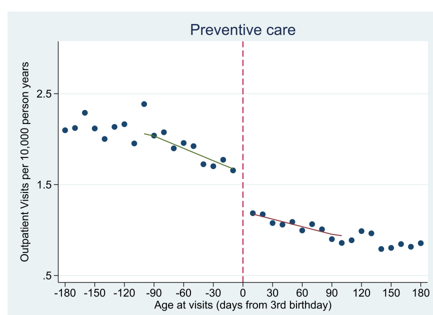

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday (90 days bandwidth). The dependent variables are outpatient visits per 10,000 person years (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

Figure 3.6: Age Profile of Inpatient Expenditure and Visits

(a) Inpatient expenditures per 10,000 person-years: 2005–2008

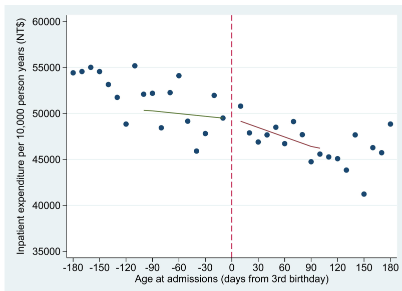

(b) Inpatient expenditures per 10,000 person-years: 1997–2001

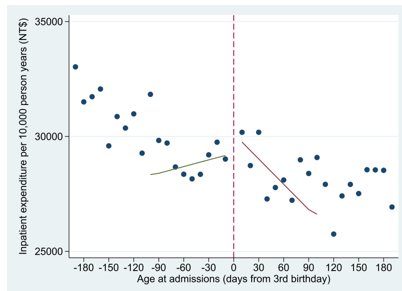

(c) Inpatient admissions per 10,000 person-years: 2005–2008

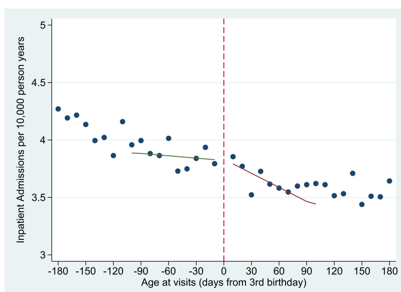

(d) Inpatient admissions per 10,000 person-years: 1997–2001

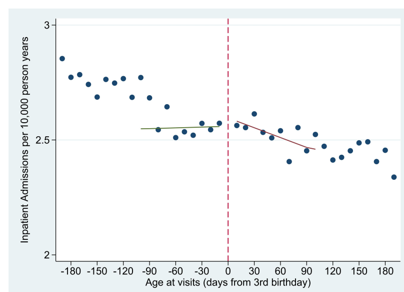

(e) Inpatient expenditures per admission: 2005–2008

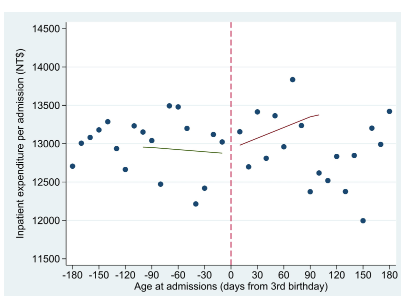

(f) Inpatient expenditures per admission: 1997–2001

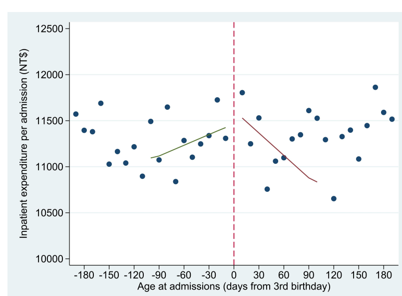

Notes: The line is from fitting a linear regression on age variables fully interacted with $Age3_i$, a dummy indicating after the 3rd birthday (90 days bandwidth). The dependent variables are inpatient expenditure per 10,000 person years, inpatient admissions per 10,000 person years, and inpatient expenditure per visit by patient's age at visit (measured in days, 180 days before and after the 3rd birthday). Each dot represents the 10-day average of the dependent variable.

3.8 Tables

Table 3.1: Patient Cost-Sharing in Taiwan NHI

	Patient Cost-Sharing			
	Major Teaching Hospital	Minor Teaching Hospital	Community Hospital	Clinic
Panel A: Outpatient service				
Copayment	360	240	80	50
Registration Fee	150	100	100	50
Panel B: Inpatient service				
1-30 days		10%		
31-60 days		20%		
after 61 days		30%		

1 USD is 32.5 NTD in 2006. For outpatient service, patient cost-sharing is through copyment. A patient pays copayment plus registration fee for each visit. If a physician prescribes a drug at a visit and the drug cost is above 100 NTD, the patient also needs to pay a share of the cost of the prescription drug, which is 20% of total drug cost. However, most visits for the children under age 3 have drug costs below 100 NTD so patients usually do not pay for their prescription drugs. On average, The out-of-pocket cost of prescription drugs per visit is very small (i.e. only 2.5 NTD). Information about copayment is from National Health Insurance Research Database codebook (2012). NHI implemented this fee schedule since July 2005. Since our sample period is from January 1st 2005 to December 31st 2008, most of outpatient visits in our sample, except visits on January 1st 2005 to June 30th 2005, are based on the above fee schedule. Before July 1st 2005, copayment for outpatient service is according to the following fee scheme: 210 NTD for major teaching hospital, 140 NTD for minor teaching hospital, 50 NTD for community hospital, and 50 NTD for clinic. Information about registration fee is from an online database of NHI registration fee survey: `http://www.nhi.gov.tw/amountinfoweb/Search.aspx?Q5C1_ID=2&Q5C2_ID=900002&Hosp_ID=1131100010&rtype=2` For inpatient care, patient cost-sharing takes place through coinsurance. Depending on the days of stay and the type of admission (acute or chronic admission), a patient is required to pay 10% to 30% of the total medical expense per admission. The above fee schedule is only for acute admission since we eliminate all chronic admissions, which only accounts for 0.3% of inpatient admissions.

Table 3.2: Weighted Average Out-of-Pocket Cost per Visit/Admission

	Out-of-pocket cost per visit/admission	
	Before 3rd birthday	After 3rd birthday
Type of Service		
Outpatient service	58.9	132.7
Inpatient service	0	1296

Note: Data are pooled NHI claims records 2005–2008. Weighted average out-of-pocket costs per visit/admission are reported in New Taiwan Dollar (NTD). 1 USD is 32.5 NTD in 2006.

Table 3.3: Selected Characteristics at Age Three before and after Sample Selection

	(1) Original Sample	(2) Continuous enrollment at age two and three	(3) Eliminating cost-sharing waiver
Male	0.525	0.525	0.524
Birith year: 2003	0.510	0.509	0.509
Birith year: 2004	0.490	0.491	0.491
1st birth	0.519	0.520	0.520
2nd birth	0.368	0.370	0.370
3rd birth (above)	0.113	0.112	0.110
Number of siblings	1.761	1.760	1.759
	(0.671)	(0.671)	(0.669)
Number of children	435,206	426,587	410,517

Note: Column (1) presents the selected characteristics for original sample: all NHI enrollees born in 2003 and 2004. Column (2) restricts the sample to enrollees who continuously register in NHI at age 2 and 3. Column (3) eliminates observations with cosh-sharing waiver, such as children with catastrophic illness (e.g. cancer) and children from very low income families since these children do not experience any price change when turning three.

Table 3.4: Descriptive Statistics

	Outpatient Service		Inpatient Service	
	Before 3rd birthday	After 3rd birthday	Before 3rd birthday	After 3rd birthday
Utilization				
Average annual visits	19.8	19.0	0.14	0.13
Average out-of-pocket cost per visit (NTD)	58.9	123.1	0	1289.7
Average medical expenditure per visit (NTD)	443.5	438.7	12980.6	13013.9
Choice of providers				
Major Teaching Hospital	4.1%	2.3%	28.4%	29.8%
Minor Teaching Hospital	5.6%	3.7%	58.6%	58.2%
Community Hosptial	3.8%	4.6%	12.8%	11.9%
Clinic	86.5%	89.4%	0%	0%
Number of children (visits $>$ 0)	375,493	364,075	13,252	12,666
Number of children-visit	2,003,097	1,954,591	19,356	18,163

Note: Data are pooled NHI claims records 2005–2008. The above descriptive statistics are based on records about outpatient(inpatient) service happened within 90 days before the 3rd birthday and 90 days after the 3rd birthday. Average annual visits is calculated by average visits at each age (measured in day) times 365. Average out-of-pocket costs and medical expenditures are reported in New Taiwan Dollar (NTD). 1 USD is 32.5 NTD in 2006.

Table 3.5: Descriptive Statistics: By Healthcare Provider

Provider	Major Teaching Hospital	Minor Teaching Hospital	Community Hospital	Clinic
Panel A:				
Composition of Medical Expenses				
Average medical expenses per visit (NT$)	986.9	697.8	528.2	338.4
Average out-of-pocket expenses per visit (NT$)	280.4	196.3	148.0	75.3
Portion of out-of-pocket expenses	0.61	0.47	0.40	0.26
Average drug fee per visit (NT$)	194.5	134.7	86.4	51.2
Average examination/treatment fee per visit (NT$)	543.6	309.9	185.4	16.7
Average diagnosis fee per visit (NT$)	204.2	206.3	213.8	255.4
Average dispensing fee per visit (NT$)	44.4	46.8	42.4	14.9
Average drug day per visit (NT$)	6.9	5.2	3.9	3.1
Panel B:				
Composition of Patients				
Male	0.54	0.56	0.56	0.54
Household income	53,485.7	47,267.9	45,750.2	47,126.2
Number of children-visit	125,847	181,767	165,695	3,483,876

Note: 1 US$ is 32.5 NT$ in 2006. The sample are all NHI enrollee born in 2003 and 2004 and continuously registered in NHI at age 2 and 3. The above descriptive statistics are based on NHI claims records within 90 days before and after enrollee's 3rd birthday. Thus, data are pooled NHI claims records 2005-2008.

Table 3.6: RD Estimates on Outpatient Care at Age 3

	2005-2008		1997-2001	
	(1)	(2)	(3)	(4)
Specification	Nonparametric	Parametric	Nonparametric	Parametric
	Local linear	Cubic spline	Local linear	Cubic spline
Visits rate at age 2	542		568	
(per 10,000 person-years)				
Bandwidth (days)	90	365	90	365
Panel A: Log(outpatient expenditures)				
*Age*3 (X100)	-6.90***	-6.99***	0.09	0.29
	[0.49]	[0.46]	[0.24]	[0.22]
Panel B: Log(number of visits)				
*Age*3 (X100)	-4.73***	-4.77***	0.22	0.20
	[0.31]	[0.32]	[0.17]	[0.16]
Panel C: Log(outpatient expenditures per visit)				
*Age*3 (X100)	-2.17***	-2.22***	-0.12	0.09
	[0.29]	[0.27]	[0.13]	[0.13]

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure, the log of the total number of outpatient visits, and the log of outpatient expenditure per visit, at each day of age. Odd columns use data within 90 days before and after the 3rd birthday (bandwidth is 90 days) and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). even columns present estimated regression discontinuities by using all available data (365 days before and after the 3rd birthday) and flexible polynominal regression (cubic spline), allowing a different slope on either side of the 3rd birthday. In the last two columns, we use the same selection criteria to create a pre-reform sample: enrollees born between 1995 and 1997 (when they are age 2 and 3). Therefore, we use 1997–2001 NHI data to obtain the above estimated results. All coefficients on *Age*3 and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 3.7: RD Estimates on Outpatient Care at Age 3: By Choice of Providers

Providers	(1) Major teaching hospital	(2) Minor teaching hospital	(3) Community hospital	(4) Clinic
Visits rate at age 2 (per 10,000 person-years)	22	30	20	469
Panel A: Log(outpatient expenditures)				
*Age*3 (X100)	-39.29***	-38.89***	17.76***	-1.92***
	[2.63]	[2.40]	[1.64]	[0.33]
Panel B: Log(number of visits)				
*Age*3 (X100)	-59.29***	-43.89***	17.71***	-1.73***
	[1.96]	[1.65]	[1.64]	[0.32]
Panel C: Log(outpatient expenditures per visit)				
*Age*3 (X100)	19.85***	5.76***	0.05	-0.19*
	[2.24]	[1.77]	[1.67]	[0.10]

Note: We collapse the individual-level data into age cells. Age is measured in days. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure, the log of the total number of outpatient visits, and the log of outpatient expenditure per visit, at each day of age. Column (1)–(4) present RD estimates of each outcome for four types of health provides by using data within 90 days before and after the 3rd birthday (bandwidth is 90 days) and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). All coefficients on *Age*3 and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 3.8: RD Estimates on Outpatient Care at Age 3: By Diagnoses, Birth Order, Gender, and Household Income

	(1) Visits rate at age 2 (per 10,000 person-years)	(2) Log(outpatient expenditure)	(3) Expenditure Elasticity
Panel A: By visit diagnoses			
URI	119	-2.38***	-0.037***
		[0.65]	[0.010]
Acute bronchitis	51	-5.56***	-0.084***
		[0.73]	[0.014]
Acute sinusitis	48	-4.10***	-0.064***
		[1.10]	[0.019]
Skin diseases	20	-14.88***	-0.259***
		[1.55]	[0.041]
Mental disorder	4	-23.18***	-0.328***
		[3.62]	[0.061]
Preventive care	2	-24.54***	-0.689***
		[6.07]	[0.29]
Panel B: By birth order			
1st birth	535	-5.97***	-0.084***
		[0.57]	[0.009]
2nd birth (above)	549	-7.90***	-0.115***
		[0.40]	[0.012]
Panel C: By gender			
Male	570	-7.65***	-0.109***
		[0.59]	[0.010]
Female	511	-5.93***	-0.085***
		[0.67]	[0.011]
Panel D: By household income			
Low	525	-6.98***	-0.101***
		[0.63]	[0.010]
High	562	-6.81***	-0.097***
		[0.54]	[0.011]

Note: We collapse the individual-level data into age cells. Age is measured in days. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure. Panel A to D report RD estimates of each outcome for various subgroups. Low income household in Panel D is defined as monthly household income is below 40,000 NTD. High income refers households with monthly household above 40,001 NTD. We use data within 90 days before and after the 3rd birthday (bandwidth is 90 days) and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday. All coefficients on $Age3$ and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 3.9: RD Estimates on Inpatient Care at Age 3

	2005–2008		1997–2001	
	(1)	(2)	(3)	(4)
Specification	Nonparametric Local linear	Parametric Cubic spline	Nonparametric Local linear	Parametric Cubic spline
Visits rate at age 2	3.9		2.5	
(per 10,000 person-years)				
Bandwidth (days)	90	365	90	365
Panel A: Log(inpatient expenditure)				
*Age*3 (X100)	-0.89	0.46	1.36	2.72
	[4.85]	[4.31]	[2.38]	[2.20]
Panel B: Log(number of admission)				
*Age*3 (X100)	-0.18	-1.26	1.14	3.12
	[2.82]	[2.56]	[2.89]	[3.13]
Panel C: Log(inpatient expenditure per admission)				
*Age*3 (X100)	-0.71	1.72	0.20	-0.40
	[3.49]	[3.21]	[2.36]	[2.48]

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents inpatient expenditures and admissions from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total inpatient expenditure, the log of the total number of inpatient admission, and the log of inpatient expenditure per visit, at each day of age. odd columns use data within 90 days before and after the 3rd birthday (bandwidth is 90 days) and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). even columns present estimated regression discontinuities by using all available data (365 days before and after the 3rd birthday) and flexible polynomial regression (cubic spline), allowing a different slope on either side of the 3rd birthday. In the last two columns, we use the same selection criteria to create pre-reform sample: enrolees born between 1995 and 1997 (when they are age 2 and 3). Therefore, we use 1997–2001 NHI data to obtain the above estimated results. All coefficients on *Age*3 and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Chapter 4

The Effect of Workplace Pensions on Household Saving: Evidence from a Natural Experiment in Taiwan

4.1 Introduction

There are several reasons why individuals may not have enough retirement saving. For example, self-control problems make people consume too much today and have too little savings for their future consumption. A recent survey shows that most Americans believe they should save more for retirement(Adams, 2014). Many countries provide public pensions to assist people to have enough retirement saving. However, population aging results in a fiscal strain on pay-as-you-go public pension systems. The use of mandatory workplace pensions is becoming a popular way for governments to increase the provision of pensions without incurring much new public spending. Several developed countries have begun to complement or substitute for public pensions by mandating workplace pensions. For example, Australia and the Netherlands have long traditions of legislation on compulsory workplace pensions.[126] This ensures that each worker is covered by a

[126] Australia introduced a new compulsory occupational pension system, Superannuation Guarantee, in 1992 that requires employers to contribute a percentage of an employee's

pension plan and therefore safeguards retirees' standard of living (OECD, 2012). In order to raise the replacement rate of pension income and mitigate the fiscal burden on the public pension system, the UK government has required all employers to provide workers with a workplace pension plan, so-called automatic enrolment, since 2012. The employers are also obliged to make employer pension contributions.

However, the ability of mandatory workplace pension schemes to raise employees' retirement wealth depends on the elasticity of substitution between workplace pensions and individual voluntary savings. If workplace pensions offset personal saving partially, such interventions could increase workers' retirement savings. If, on the other hand, workplace pensions substitute perfectly for private saving, then legislation requiring employers to offer pensions for their workers may just generate a deadweight loss and fail to help employees accumulate more wealth for their retirement.[127]

In this paper, I estimate the causal effect of workplace pensions on household saving by analyzing a pension reform in Taiwan which mandated, from 2005 onwards, that all private-sector employers should pay a minimum contribution of 6% of each employee's wage, to the latter's individual pension account, monthly. Before the reform, most private-sector employees in Taiwan did not obtain employer-sponsored pensions when they retired. Thus, this reform has substantially increased the pension coverage of private-sector workers and raised employers' pension contributions. I exploit this policy change to obtain exogenous variation in employer's pension contribution for the affected workers and employ a difference-in-differences approach to overcome the potential endogeneity problems when estimating the effect of workplace pensions on household saving (Gale, 1998). I find the pension reform significantly reduces the household saving rate (as a percentage of disposable income) of private sector employees by 2.06 percentage points to 2.45 per-

salary into the latter's individual pension account (Tapia, 2008). The Netherlands also has mandatory occupational pensions covering more than 95% of employees (Tapia, 2008).

[127]For example, this deadweight loss may come from the administrative cost of implementing pension law or distortions of the labor market. If employers can fully shift pension costs to workers by reducing employees' wages, then there would be no distortion in the labor market.

centage points, suggesting the elasticity of substitution between workplace pensions and household saving is about -0.50 to -0.60.[128] Since workplace pensions do not crowd out household saving completely, my results suggest that making workplace pensions compulsory might be an effective policy to raise workers' retirement wealth.

This paper contributes to the current literature in two important dimensions. Firstly, to the best of my knowledge, this paper is the first study to use a national policy change as a natural experiment to identify the causal effect of workplace pensions on voluntary household saving of affected workers. Many early studies used ordinary least squares (OLS) regression to estimate the offsetting effect of workplace pensions on saving and their results were mixed. Many of them suggest that workplace pensions have a very small and insignificant effect on household saving (Cagan, 1965; Katona, 1965; Hemming and Harvey, 1983; Hubbard, 1986; Gustman and Steinmeier, 1999; Alessie et al., 1997). In contrast, a few studies find that workplace pensions may substantially crowd out 49% to 92% of other household savings (Munnell, 1976; Gale, 1998; Euwals, 2000). As Gale (1998) points out, the magnitude of OLS estimates may be upwardly biased towards zero since the estimated offsetting effects are confounded with unobserved heterogeneity in saving preferences. For example, employees with a strong propensity to save for retirement may choose jobs offering generous pension plans. This unobserved preference heterogeneity would introduce a positive correlation between workplace pension wealth and household savings. Hence, the OLS-estimated savings-offsetting effects of workplace pensions will tend to be underestimated.[129] To obtain unbiased estimates of offset effect, it is necessary to find exogenous variation in workplace pension wealth that may be driven by an exogenous policy change or instrumental variables (IV). To my knowledge, only a recent study by Engelhardt and Kumar (2011) tries to solve this endogeneity problem.[130] They use US employer-provided pen-

[128] Namely, my results imply that a 10% increase in employer pension contributions rate is offset by a 5–6% reduction in household saving rate.

[129] In terms of absolute value.

[130] Chetty et al. (2014) use Danish administrative data to investigate the effect of workplace pensions on employees' saving behavior. They utilize the variation in employer

sions Summary Plan Descriptions, legal descriptions of pensions, matched to Health and Retirement Survey (HRS) respondents, and then employ this detailed information on pension rules to construct their instruments. Their OLS results reveal that workplace pension wealth has no effect on non-pension wealth, but their IV estimates show that workplace pensions offset 53–67% of household savings, which is quite similar to my results. My methodology is different from theirs. I exploit a reform-induced increase in workplace pension contribution for private-sector employees in Taiwan and use an unaffected sector, namely civil servants and national enterprise workers, as a comparison group to control for other, unobserved confounding effects. By using this difference-in-differences framework, I am able to acquire causal estimates of the effect of workplace pensions on employees' saving.

Secondly, this paper exploits a policy change that is more transparent and easier to interpret theoretically compared to prior literature. Several recent studies (Attanasio and Brugiavini, 2003 ; Attanasio and Rohwedder, 2003; Feng et al., 2011; Banerjee, 2011) form their instruments for public pension wealth by exploiting changes in the rules on public pension benefits, for example, various increases in the public pension benefit eligibility age for different cohorts or changes to the indexation of the benefit, to generate convincingly exogenous variation in public pension wealth across different cohorts and occupational groups and then estimate the offsetting effect of public pensions on non-pension saving. However, the policy changes used in these prior studies are usually not uniform across households, can be hard to characterize, and might depend on unknown parameters such as the discount rate. By contrast, the reform used in this paper uniformly

pension contributions when employees switch jobs and find that workplace pensions only offset 10%–15% of other household savings. Although they performed many robustness checks on their results, their crowding-out estimates may still be downwardly biased (in absolute value) since job switches are endogenous and the variation in employers' pensions contributions induced by firm switches may be correlated with employees' savings preferences. In addition, the reform used in this paper affects most private sector workers in Taiwan. If there exists adjustment cost in saving decision, such big changes in pension coverage is more likely to make workers re-optimize their saving decisions, which could also explain why I find larger crowding-out effect than the estimates in Chetty et al. (2014).

imposes that 6% of wage income be invested in illiquid instruments. This feature allows the paper to contribute a clean estimate of substitutability between illiquid pension savings and household savings, obtained with a transparent methodology. In addition, private-sector workers comprise 85% of employees and more than 60% of the labor force in Taiwan and the pension reform raised the coverage rate of workplace pensions for private-sector workers from 44% to 100% in a very short amount of time (Taiwanese Labor Statistics, 2007). Figure 4.1 reveals a salient change-up to a 56% increase-in pension coverage for private-sector employees after the reform.[131] This sudden expansion in workplace pension coverage in Taiwan gives us a rare chance to estimate closely the average treatment effect (ATE) of workplace pensions on employees' saving behavior.

The rest of the paper is organized as follows. Section 4.2 gives a brief overview of the employer pension system in Taiwan and introduces the context of the mandatory workplace pension reform implemented in 2005. Section 4.3 introduces the data and defines the treatment and control groups. Section 4.4 describes my empirical strategy. Section 4.5 analyzes the main results. Section 4.6 performs various specification checks. Section 4.7.1 discusses the distinct impact of the reform across the saving rate distribution. Section 4.8 provides concluding remarks.

4.2 Policy Background

Prior to the 2005 pension reform, the private-sector pension system in Taiwan was legislated by the Labor Standard Law that was established in 1984. The law required employers to make flexible pension contributions, of 2–15%

[131] This measure could be the "lower bound" of policy variation, since the pension coverage rate only indicates the percentage of employees covered by workplace pension plans. It does not mean that these covered workers will be eligible for workplace pensions when they retire. In particular, prior to the reform, the vesting period for pension benefits was very long (25 years in the same firm) so that many private-sector workers could not obtain a pension even if their company offered a pension plan. The 2005 pension reform introduced immediate vesting and made all private-sector workers eligible for pension benefits after retirement. Therefore, the "true" increase in pension coverage induced by the reform could be even larger than 56%. I will discuss this issue in Section 4.4 and Section 4.5.

of an employee's wage, to retirement funds "owned by firms." However, the vesting period was very long. The employees had to stay at the same firm for 25 years, or for 15 years until they were at least 55 years of age. Since the average lifespan of companies in Taiwan is 13 years and the average job tenure is only 6 years (Yang and Luoh, 2009), most private-sector employees, other than senior workers in big firms, did not expect to obtain their pensions. Only 10% of firms obeyed the law in setting up their company pension funds.

To increase the coverage of workplace pensions, the Legislative Yuan (Taiwanese Congress) approved the New Labor Pension Law in July 2004 and implemented it one year later (July 2005). The main features of the new workplace pension regulation were as follows. Firstly, the new pension scheme introduced compulsory workplace pensions. The employers were mandated to make monthly pension contributions of at least 6% of an employee's wage to the latter's individual pension account.[132] According to the 2005 Taiwanese Labor Statistics,[133] in 2004 (i.e. one year before the reform), only 20% of private-sector retirees were eligible for workplace pensions. In other words, the reform caused 80% of private sector workers to be newly eligible for workplace pensions. Secondly, the new system provided immediate vesting, that is, eligibility for pension benefits would be unrelated to a worker's current job tenure. In fact, employees' pension benefits would now accumulate in their personal accounts rather than in their firms' pension funds. Consequently, under the new pension system, employees would be assured that they would receive their pensions as long as their employers' paid the monthly contribution.[134] Table 4.1 briefly compares the new and old workplace pension systems in Taiwan.[135]

The new pension scheme applies automatically to all private-sector work-

[132]Self-employed workers were not affected by reform. Therefore, my sample does not include self-employed workers.

[133]These are published by the Taiwanese Council of Labor Affairs.

[134]However, workers can only use the money in this account after they retire and the retirement age is 60. Before retirement, the government helps employees to invest their pensions and guarantees a minimum rate of return.

[135]The pension contribution is not taxable until retirement. When people start to receive their workplace pensions, this income would be taxed.

ers joining the labor market for the first time and to those who switch jobs after the reform. Employees who stayed in their current job after the reform were given the option of either staying on their old pension scheme or changing to the new pension plan, within a transition period of five years. At the end of the transition period, all employees had to select their workplace pension plan. The old pension scheme had a higher income replacement rate than the new scheme when workers became eligible for pension benefits. In addition, employees who switched to the new system had to give up all benefits amassed in the old pension system.[136] I expect that only senior workers who were close to retirement and had accumulated substantial pension wealth in their company's pension fund would have been likely to have continued with the old pension plan and thus been unaffected by the pension reform. As reported in the 2006 Taiwanese Labor Statistics, the coverage rate of the new pension scheme decreases with workers' age. The coverage rate for employees under 50 years of age is 84%, but for workers over 50 it drops to 48%.[137]

In contrast to private-sector workers, public-sector employees, including civil servants and workers in national enterprises, have their own pension systems that were not affected by the 2005 pension reform.[138] Taking advantage of this institutional difference between the two sectors, I use public-sector workers as the control group to identify the causal effect of workplace pension provisions on household savings. In Section 4.5, I will examine the validity of the control group.

[136] Since workers changing their jobs in the future would not have been able to obtain these pension benefits under the old pension system, many employees are likely to have switched to the new system even if they had accumulated some pension wealth in the old system.

[137] In the old pension system, there was a maximum tenure of pension contribution (30 years). After that, employees have to switch to the new pension system.

[138] In fact, the public-sector pension system was not changed at all during my sample period (2002–2008).

4.3 Data and Sample

4.3.1 Data

In order to calculate the household saving rate and identify the targeted sample, data recording detailed information on household income, consumption, and household members' occupations are required. I use the Taiwanese Survey of Family Income and Expenditure (TSFIE), conducted annually since 1976 by the Taiwanese Directorate-General of Budget, Accounting and Statistics (DGBAS). The TSFIE is an ongoing repeated cross-section income and consumption data set and provides a nationwide representative sample of Taiwanese households. Its sample size is around 14,000 households and 55,000 individuals each year. I use detailed information on household income source and expenditure in TSFIE to calculate household saving rate. Also, TSFIE includes information on household members' occupations and working status. I use this information to define the treatment and control groups.

4.3.2 Sample

I employ six years of TSFIE data from 2002 through 2008 (excluding the reform year of 2005). Since the new workplace pension system was introduced in 2005, I use the 2002–2004 and 2006–2008 samples to represent the periods before and after the pension reform, respectively. I confine the sample to families headed by prime-age workers (20–50 years old) for two reasons.[139] First, the main purpose of this paper is to investigate whether mandatory workplace pension provision can raise employees' retirement "wealth." Since retirement wealth is comprised of prime-age saving, and the average retirement age of private-sector workers in Taiwan is around 55 (Taiwanese Labor Statistics, 2005), it is better to focus on pre-retirement (prime-age) saving behavior rather than that exhibited in old age. Secondly, the coverage rate

[139] Since I focus on employees' saving behaviors, I exclude households whose wage income is zero. I also exclude households whose family members are self-employed or work in the agricultural sector, since workers from these sectors often misreport their income and consumption (Gale, 1998; Attanasio and Rohwedder, 2003).

of the new pension system decreases with workers' age; most employees under 50 years of age are covered by the new pension scheme. In contrast, more than half of all workers above 50 stayed in the old pension system. Hence, this sample selection also ensures that most of those in our treatment group really have been affected by the pension reform. Thus, the original sample size for the prime-age-employee households is 35,775. To avoid the effect of outliers, I exclude households whose saving rate is above 100% or below −100%. This reduces the sample size to 35,715.

In contrast to previous studies (Aguila, 2011; Chou et al., 2003), which define the treatment group by the sector of the head of the household, I identify the treatment group using detailed sectoral information on household members drawn from the TSFIE. Hence, the treatment group consists of those households with at least one member working in the private sector and no one working in the public sector. In the same way, the control group contains those with at least one member working in the public sector and no one working in the private sector. This arrangement is more suitable for Asian families, since the family size tends to be larger, with many family members participating in the labor market.[140] If I did not take each family member's employment sector into account, the estimates of the policy impact would be biased toward zero when the head of the family and other family members were working in different sectors.[141] After applying this requirement, the final sample includes 32,869 households, of which about 28,729 (87.31%) belong to the treatment group (private-sector families) and 4,140 (12.69%) to the control group (public-sector families).

[140] In my sample period (2002–2008), the average number of family members is 3.88 and more than one family member (1.69 people) has a job.

[141] For example, suppose the head of the household is a public-sector worker and other family members are private-sector workers. Using the head of the household's sector as the criterion, this family would be defined as part of the control group. However, this household's savings should be affected by the pension reform since other family members are in the treatment group. These contaminated households are excluded from my sample. I allow only purely private (purely public) sector families in my sample. I will conduct a robustness check on this issue in Section 4.6.

4.4 Empirical Specification

In this section, I estimate a difference-in-differences model comparing the evolution of the saving rates in private-sector and public-sector households around the time of the pension reform. This strategy will identify the impact of the mandatory pension reform as long as there were no other reasons for a change in the relative saving behaviors of private- and public-sector employees at that time. The following difference-in-differences regression is used for my main analysis:

$$SR_i = \beta^{DD} PENSION_i + \alpha PRIVATE_i + \gamma YEAR_i + X_i \psi + \varepsilon_i \tag{4.1}$$

Where the outcome variable of interest, the household saving rate SR_i is measured as the difference between household disposable income and consumption expenditure divided by household disposable income. The household total income includes wage income and non-wage income (e.g. non-wage benefit, asset returns, and transfer income). I subtract any income tax, capital tax, and employee mandates for health insurance from household income to get household disposable income. The consumption expenditures includes spending on both durable and non-durable goods. Using the saving rate as the main outcome variable can help me directly calculate the elasticity of substitution between workplace pension and household saving. In particular, I do this by comparing the employer's pension contribution rate and the decreased size of the household saving rate induced by reform. I also use level of household savings (take log) as the outcome variable for robustness check.

I include $PRIVATE_i$, a private sector dummy, where $PRIVATE_i = 1$ represents the treatment group (private sector households) and $PRIVATE_i = 0$ denotes the control group (public sector households). $YEAR_i$ are year dummies for each year in the sample period: 2002 to 2008 (except for the reform year 2005). The parameter α measures unobservable time-invariant differences in saving rates between private- and public-sector households.

The parameter γ captures year-fixed effects (common macroeconomic impacts).

Since differences in observed covariates may lead to distinct time trends in the household saving rate between the treatment and control groups, it is necessary to control the related covariates to eliminate the impact of these other confounding effects. Controlling covariates can also reduce the residual variance of the regression and produce more efficient estimates (Andrietti and Hildebrand, 2006; Meyer, 1995). I include a number of covariates X_i that could affect household saving suggested by previous studies (Aguila, 2011; Chou et al., 2003): (1) Family characteristics: head of family's age, age squared, their education and their gender; spouse's education; number of children under 18, number of household members over 65, number of members aged 18–64, number of working members, and dummies for area of residence. (2) Industry and occupation: head of family's industry and occupation. (3) Household wealth: household non-wage income; housing assets.[142] ε_i is the error term.

The key variable $PENSION_i$ is a dummy indicating that household i is affected by pension reform, meaning that household i has someone working in the private sector in the post-reform years 2006–2008.[143] Its coefficient β^{DD} is the standard difference-in-differences estimator. Since I control for group and year-fixed effects, β^{DD} measures the differential trend in average household saving rates among private sector workers relative to public-sector workers in the post-reform years.

I can attribute the difference in the evolution of the household saving rate between the two groups to the effect of the mandatory pension reform on the saving rate among private-sector households if I impose the following identification assumptions: Firstly, the public- and private-sector households' saving rates would follow a common trend in the absence of the pension policy change. Given this assumption, I can use the observable post-reform trend in the saving rates among public-sector households to derive

[142] Dummy variables for owning one's house and house size.

[143] $PENSION_i = 0$ means either that a member of household i works in the public sector or that the observation is from the years 2002–2004.

the counterfactual evolution of private-sector household saving rates after the reform. This assumption will ensure that my results do not come from different pre-reform trends in household saving rates between the treatment and control groups.

Secondly, except for the pension reform, no other shock during my sample period has a differential effect on the two groups' household saving rates. That is, the difference in the evolution of the saving rates for private- and public-sector households after 2005 should be driven by the pension reform. I cannot completely rule out that some other shock during this period might have had distinct effects on private- and public-sector households' saving. However, given the magnitude of the pension coverage change,[144] it seems highly unlikely that other shocks could be the driving force for the relative shift in saving rates between the two sectors of workers over this time period.[145]

Thirdly, I assume there is no selection into the treatment based on unobservables ε_i once I control for the observable covariates X_i. In other words, I assume that employees with a high savings preference (unobserved) do not switch to the private sector to obtain workplace pensions after the reform. This assumption ensures that my results are not driven simply by workers' self-selection into the treatment group after reform. Such a self-selection problem is unlikely to occur, since public-sector employers offer more generous pensions to their employees than private-sector employers even after the pension reform and the substantial increase it brought about in most private-sector workers' workplace pension wealth.[146] In Section 4.5, I will use three placebo tests to examine the credibility of these identification as-

[144] As mentioned in Section 4.1, this reform made around 80% of private-sector employees newly eligible for a workplace pension when they retire.

[145] In fact, during the entire sample period (2002–2008), no other policy change was made affecting the labor market or workers' savings behavior. Thus, it is arguably sound to make this assumption.

[146] Since people need to pass exam to become public sector workers, I find the pass rate of exams did not change a lot during my sample period. The pass rate was 4.9% in 2004 and then became 5.5% in 2007. In general, public positions are highly selective. There is considerable excess demand for these positions even after the pension reform. In addition, such self-selection may lead to my estimates being underestimated.

sumptions.

Two caveats to my estimating procedure have to be mentioned before I analyze the results. First, my identification strategy indeed analyzes the intention-to-treat effect instead of the average treatment effect on the treated. Since I do not have individual-level data concerning workers' workplace pension coverage, I estimate the reduced-form effects on all private-sector workers (i.e. the eligible population) rather than on workers who are newly covered by workplace pensions (i.e. the affected population). To recover the average treatment effect on the treated, I need to divide β^{DD} by the proportion of truly affected private-sector employees (Baker et al., 2008; Bloom, 1984).[147] If the fraction of truly affected workers is close to one, the intention-to-treat effect will approach the average treatment effect on those who are treated. In my sample (prime-age workers), over 84% of employees are covered by the new pension scheme. Furthermore, this reform has likely made 80% of private-sector retirees newly eligible for workplace pensions. It may be argued that the fraction of affected employees is high. Hence, β^{DD} (the intention-to-treat effect estimate) may not be far from our quantity of interest, namely, the average treatment effect on the treated. I will return to the specifics of this issue in Section 4.5.

Second, the correct computation of standard errors is a crucial issue in the difference-in-differences approach. Since the policy variation I use is at the sector-year level, I present the standard errors clustered on the sector-year cells to account for any dependence of the unobservable error within sector-year cells.[148] Furthermore, in recognition of the small number of clusters, following Cameron et al. (2008)'s suggestion, I adopt relatively conservative inference by using the $T(G-2)$ distribution rather than the standard normal distribution to form critical value and p-values.[149] For my small number of clusters, this correction may make a substantial difference to the inference results (Angrist and Pischke, 2009). In Section 4.6, I will

[147] It also needs to be assumed that there are no externalities of the pension reform for treated households.
[148] There are two sectors (private and public) and six years (2002–2004 and 2006–2008). Therefore, I have $2 \times 6 = 12$ clusters.
[149] G is the number of clusters, namely, 12.

also use different levels of clustering (e.g. clustering by sector and by pre- and post-reform periods) and various inference methods, such as block bootstrap (Bertrand et al., 2004) and wild bootstrap (Cameron et al., 2008), to address this issue.

4.5 Results

4.5.1 Summary Statistics

Table 4.2 compares the trends in the outcome variables and covariates between the treatment and control groups before and after the reform. It also presents the results of simple difference-in-differences estimates for each variable.[150] Two things can be learnt from Table 4.2. First, for private-sector workers, the household saving rate, household savings and wage income decrease significantly after the reform.[151] However, the corresponding variables for public-sector households remain the same after the reform. In fact, the average growth rate of per capita GDP in Taiwan during 2006 to 2008 is around 4%. Therefore, one might expect wage income as well as household savings not to decrease in the absence of pension reform. The simple difference-in-differences estimates indicate that the pension reform reduced the household saving rate, household savings and wage income of private-sector workers significantly by 1.62%, 34,022 NT$, and 36,520 NT$, respectively (all annual numbers).[152] The finding of decrease in wage income indicates that pension reform leads to wage income being reduced by 5.31% and employers may fully shift the mandatory pension cost to the employ-

[150] The simple DD estimates I employ here are

$$Variable_i = \delta^{DD} PENSION_i + \delta_1 PRIVATE_i + \delta_2 POST_i + \varsigma_i$$

where $PENSION_i$ and $PRIVATE_i$ are as defined in Section 4.4. $POST_i$ is the dummy for the post-reform period: 2006–2008. I focus on the DD estimates δ^{DD}.

[151] All of these variables are measured in 2007 New Taiwan Dollars.

[152] 1 US$ = 32.4 NT$ in 2007.

ees,[153] which is consistent with the findings in Yang and Luoh (2009).[154] Thus, the pension reform may have little impact on the lifetime income of private-sector households (i.e. no income effect on saving).

Second, except for the above three variables and dummies for area of residence, the simple difference-in-differences estimates reveal that other variables exhibit different trends between private- and public-sector employee households after the pension reform. In other words, there is little evidence of a composition change in the covariates after the reform. I will control these covariates further in the following regression analysis.

4.5.2 Main Results

Columns (1)–(4) of Table 4.3 report the estimated coefficients on the key variable $PENSION$ in difference-in-differences estimation (equation (4.1)). I begin by presenting the estimate from the basic difference-in-differences regression controlling for a dummy for private-sector households and year fixed effects (Column (1)). Then, I gradually include the covariates of family characteristics, working industry, occupation and household wealth (Column (2)–(4)). The fact that the estimates are quite stable across different specifications is comforting. All of the estimates are significantly different from zero at the 5% level. My preferred specification (column (4)) implies that pension reform causes the household saving rates of private sector employees to fall by 2.06 percentage points. This is a sizeable decrease that amounts to around 10% of the pre-reform household saving rate.[155]

Note that this estimate is an intention-to-treat effect. To arrive at the average treatment effect on those treated, it must be divided by the proportion of truly treated workers. I propose two measures of this proportion. Firstly, as reported by the 2006 Taiwanese Labor Statistics, around 84% of workers below age 50 are known to be covered by the new scheme. Using

[153] This is a simple calculation: $36,520/673,430 = 5.31\%$, $673,430$ is pre-reform wage income.

[154] Yang and Luoh (2009) use 2003–2007 Manpower Utilization Survey (i.e. labor force survey data, like Current Population Survey in the United States) and find pension reform reduces the hourly wage rate by 5.92%.

[155] The pre-reform household saving rate of private sector workers is 22%.

this number to represent the probability of treatment for my sample, the resulting estimated average treatment effect on the treated indicates that the pension reform leads to a 2.45 percentage points decline in the household saving rate of private-sector workers.[156] Secondly, I use the change in the coverage rate of workplace pensions to get an estimate of the percentage of employees treated. As mentioned in Section 4.1, Figure 4.1 indicates that the reform induces a 56% increase in the coverage rate of workplace pensions. The estimate of the average treatment effect on the treated calculated by this method suggests that making workplace pensions compulsory reduces the household saving rate of private-sector workers by 3.68 percentage points.[157] I prefer the first estimate to the second since, with the long vesting period, many covered employees still failed to qualify for a pension when they retired under the old scheme. In fact, before the reform, only 20% of retired private-sector workers were eligible for workplace pensions. The estimate of the proportion of truly treated workers derived using the second method could substantially underestimate the probability of treatment.

These estimates of the impact of the reform on the household saving rate can be used to estimate the elasticity of substitution between workplace pensions and household savings, which is a major issue in previous studies. The statutory contribution rate of employers is at least 6% of a worker's wage[158] and the average labor income share for private sector households in my sample is around 68%.[159] The simple calculation implies that the contribution rate of workplace pension is around 4.08% of household disposable income. Comparing the estimates of the reform impact on household saving rate with pension contribution rate, I obtain an implied degree of substitution between a workplace pension and households' voluntary saving of around -0.50 to -0.60.[160] That is, a one dollar increase in the workplace pension

[156]This is just a simple calculation: $2.06/0.84 = 2.45$.

[157]This is again just a simple calculation: $2.06/0.56 = 3.68$.

[158]According to the 2006 Taiwanese Labor Statistics, most employers pay only the minimum pension contribution, that is 6% of workers' wages.

[159]First, I calculate each household's wage income share (i.e. wage income divided by household disposable income). Then, I compute the mean of wage income share by using the sample of private-sector households during the pre-reform period.

[160]This is a simple calculation using the estimate of the intention-to-treat effect:

is likely to displace 50 to 60 cents of household savings.

4.5.3 Falsification Tests

In this section, I utilize three falsification tests to check the validity of public-sector households as a comparison group. The first placebo test uses previous periods (2000–2004) of TSFIE data to check whether there exists a parallel trend in saving rates between private- and public-sector households before the reform. I assign a fake policy change to 2002 and choose 2000–2001 and 2003–2004 as the pre- and post-reform periods, respectively. During this period, there was no policy change affecting workers' saving behavior in either sector. If the two groups' household savings show a common trend before the pension reform, then insignificance in the "treatment effect" ($\widehat{\beta}^{DD}$) estimates should be expected in the 2000–2004 sample. The result in Table 4.3 column (5) indicates that the point estimate of $\widehat{\beta}^{DD}$ is only -0.0053 and is not significantly different from zero. This result implies that private- and public-sector household saving rates might have shared similar trends before the 2005 pension reform.

Secondly, I use "less affected" private sector households whose heads of the family work in the banking industry as a new treatment group[161] to examine whether there are other confounding factors affecting the saving rate of private and public sector households differently. Before the pension reform, the workplace pension coverage rate in the banking industry was particularly high, with around 90% of employees having a workplace pension plan (Taiwanese Labor Statistics, 2001). The pension reform should have had a much smaller impact on these workers. Hence, if there is no shock other than the pension reform having a distinct impact on the saving rates of private- and public-sector employees, the workers in private banks should have less of a savings response to the reform. Column (6) in Table 4.3 indicates that the pension reform led to a small and insignificant reduction in the saving rate of households with heads working in private-sector banking

$-2.06/4.08 = -0.50$ and the estimate of the average treatment effect on the treated: $-2.45/4.08 = -0.60$.

[161] The comparison group is the same (public sector employees).

(the point estimate is only -0.0050).

The last placebo test uses that part of the sample with household heads on the verge of retirement (i.e. age of the head of household is above 51 years). Since the average age of retirement in Taiwan is around 55, and over half of 51–65-year-old private-sector workers are still accumulating their pension wealth under the old system, their savings behavior should be expected to show less of a response to the mandatory workplace pension reform if there was no other shock affecting private- and public-sector workers' saving behavior differently at the time. The result, shown in Table 4.3, Column (7) confirms my conjecture that the point estimate of $\widehat{\beta}^{DD}$ would not be significantly different from zero and indicates that the evolution of the saving rate for private sector employees close to retirement did not change after the reform.

In Table 4.3, I also test for equality between my main estimate (Column (4)) and the estimates in the three placebo tests. All of the placebo test estimates are significantly different from the main estimate, lending credence to my expectations regarding the impact of the reform on households' voluntary saving. In sum, these placebo tests imply that the pension reform might be the main reason for the difference between the trends in the average household saving rate for private relative to public-sector workers after 2005.

4.5.4 Magnitude of Estimates

My results reveal a substantial offsetting effect of workplace pensions on household savings, with the elasticity of workplace pensions to household saving estimated at around -0.50 to -0.60. This finding is in accordance with recent studies that have employed an IV identification strategy to estimate the crowding-out effect of public pension or workplace pensions on household savings/wealth.

Attanasio and Brugiavini (2003) explore the effect of social security reform in Italy and obtain an average elasticity of substitution between public pension and household saving of -0.35 to -0.71 (across all age cohorts: 20–

65) and find the substitutability particularly high (close to -1) for employees aged 35–45. Using UK household data, Attanasio and Rohwedder (2003) find that the estimated crowding-out effect of public pension on household savings is not significant for the young cohort (20–31 year-olds) but that the elasticity of substitution is significant and around -0.65 to -0.75 for middle-aged households (43–64 year-olds).

Engelhardt and Kumar (2011) use HRS data and construct IVs for workplace pensions by exploiting the detailed employer-provided pension plan supplement that comes with the HRS. In contrast to this paper, their HRS sample contains households whose heads are aged above 50 years old. However, it is still appropriate to compare their results with my difference-in-differences estimates since they estimate the effect of pension "wealth" on household "assets" (stock variable) rather than savings (inflow variable), and retirement wealth mainly consists of prime-age savings. Actually, they find a similar magnitude of substitution between workplace pension wealth and other household assets. That is, workplace pension wealth offsets half (53–67%) of non-pension wealth in the United States.

4.6 Specification Checks

My estimations presented thus far clearly show that the pension reform has reduced the private-sector family saving rate by between 2.06 and 2.45 percentage points. The results indicate a considerable elasticity of substitution of workplace pension for household savings, at around -0.50 to -0.60. I also conducted three falsification tests to check the effectiveness of the control group and confirmed that public-sector families probably make a suitable control group. I now experiment with various specifications to examine the robustness of my results.

4.6.1 Different Methods of Statistical Inference

In this section, I explore the robustness of my main results to alternative methods of statistical inference. Firstly, to account for temporal dependence in the error, I start by assuming that the dependence is restricted to the pre- and post-reform periods. However, I find that the standard error becomes even smaller than in my main results, where the standard error is clustered on sector-year cells. The point estimate is significantly different from zero at the 1% level when I use the $T(G-2)$ distribution to obtain the p-value (see Column (1) in Table 4.4).[162]

Next, following Bertrand et al. (2004), I conduct a block bootstrap procedure (clustered on sector-year cells). This method maintains the correlation structure within clusters by keeping samples that belong to the same cluster together in a block. Furthermore, instead of bootstrapping the standard error, this method directly bootstraps t-statistics so that I can only report the p-value and not the standard error. The result shows that the p-value computed by a block bootstrap, achieves the 1% statistical significance level (see Column (2) in Table 4.4).

Finally, I also show results from the wild cluster bootstrap approach[163] suggested in Cameron et al. (2008). This method is aimed at improving inference in cases with a small number of clusters and avoids the problem of inestimable coefficients by resampling the residuals rather than pairs of independent and dependent variables (e.g. block bootstrap). The inestimable problem is more serious in this paper since I have small clusters and the parameters of interest are indicator variables (difference-in-differences estimator $\widehat{\beta}^{DD}$), producing the problem of the resampling regressors are all being 0 or 1. The p-value computed by this approach is, as expected, a bit larger (0.007) but still indicates the 1% significance level (see Column (3) in Table 4.4).

4.6.2 Different Definition of Saving Rate

Deaton and Paxson (2000) suggest that the household saving rate can be approximated by the difference between the logarithm of a family's after-tax income $\ln(Y)$ and the logarithm of the family's consumption expenditure

[162] G is the number of clusters, making it 4 in this case.
[163] Again, I cluster on sector-year cells.

$\ln(C)$. Hence, I redefine my dependent variable, household saving rate, as $\ln(Y)-\ln(C)$. Column (4) of Table 4.4 indicates that pension reform induces a 2.50 percentage points decline in the household saving rates of private sector employees. The point estimate is less precise[164] and significantly differs from zero at the 5% level (p-value is 0.029).

4.6.3 Different Definitions of Treatment and Control Groups

In the third column (5) of Table 4.4 we follow previous studies (Aguila, 2011; Chou et al., 2003) and redefine the treatment and control groups using the head of the household's work sector. As mentioned in Section 4.3, this specification should lead to estimates of the pension crowding-out effect $\widehat{\beta}^{DD}$ that are upwardly biased toward zero when other family members have jobs in different sector. As expected, I find the estimated reform impact to be smaller, suggesting only a 1.37 percentage points decrease in the household saving rate, but is not statistically different from my main estimate.

4.6.4 Different Sample Periods

To eliminate any influence of the anticipation of the reform on my results, I also exclude 2004 TSFIE data. This is the year in which the new pension law was passed. However, the result based on this sample period is quite similar to my main result (see Column (6) in Table 4.4).

4.6.5 Controlling Household Earned Income

To get some sense of how my main results depend on any change in household earned income induced by reform, I include household earned income in my regression. After controlling household earned income, the estimated reform impact is smaller, namely, 1.38 percentage points decrease in the household saving rate. However, this estimate is not statistically different from my

[164] Standard error is 0.010.

main estimate (see Column (7) in Table 4.4).[165]

4.7 Impact Across the Saving Rate Distribution

The estimates in the previous sections show the "average" reform impact on workers' voluntary saving. I find that mandating workplace pensions for employees, on average, reduces prime-age workers' saving rate in the treatment group by between 2.06 and 2.45 percentage points. However, these estimates summarize the reform's impact in a single number and do not give us the "overall" reform impact on a worker's saving behavior if the pension reform did not have a uniform effect on each worker. This conjecture is highly possible since around 20% of private sector workers had been eligible for workplace pension before 2005 pension reform and they should be less affected by the reform. In addition, these workers could have stronger preference toward saving so they might choose the jobs that provide generous workplace pension and have higher household savings.

In this section, I explore the possible heterogeneity of the savings response to mandatory workplace pensions across the saving rate distribution. This analysis will give us a more complete picture of how mandatory workplace pensions affect workers' saving behavior. It will also provide useful lessons for other countries (e.g. the UK) implementing similar mandatory employer pension policies.

4.7.1 Quantile Differences-in-Differences Estimation

To examine how the effect of mandatory workplace pensions differs across households with different saving rates, I use a quantile difference-in-differences regression to estimate the policy effect on the "entire" distribution of private-sector household saving rates. Since the impact of the pension reform over

[165] I also run a regression of the savings rate on log household earned income to know savings rate-wage relationship. I find that 1% increase in household earned income is associated with 0.15% increase in household saving rate. If employers can fully shift cost of pension contribution by reducing 6% of worker's wage, the saving rate could be reduced by 0.9%.

the unconditional distribution of household saving rates is my outcome of interest, I adopt a recently developed estimation technique of unconditional quantile regression (Firpo et al., 2009) to obtain the quantile difference-in-differences estimator for each quantile. The traditional quantile regression method (Koenker and Bassett, 1978) cannot estimate the treatment effect on unconditional quantiles. Firpo et al. (2009) address this issue by replacing the outcome variable (i.e. household saving rate) with recentered influence function (RIF) and then conducting a standard OLS regression. The RIF in this paper is defined by:

$$RIF(SR_i, Q_S(\theta)) = Q_S(\theta) + \frac{\theta - 1\{SR_i \leq Q_S(\theta)\}}{f_{SR}(Q_S(\theta))}$$

where $Q_S(\theta)$ is the θth quantile of the household saving rate. $1\{.\}$ is an indicator function and $f_{SR}(Q_S(\theta))$ is the density of the household saving rate at the θth quantile. $\dfrac{\theta - 1\{SR_i \leq Q_S(\theta)\}}{f_{SR}(Q_S(\theta))}$ is the influence function for evaluating the effect on the estimates of the quantile producing by changing one data point in the sample. The key feature of the RIF is that the expected value of the RIF (conditional on the covariates X_i) is equal to the unconditional quantile of the saving rate $Q_S(\theta)$. Applying this property, Firpo et al. (2009) show that we can obtain the estimates of the covariates' unconditional quantile effect by simply using an OLS regression of the RIF on the covariates. I estimate the following quantile difference-in-differences regression:

$$RIF(SR_i, Q_S(\theta)) = \beta^{DD}(\theta)PENSION_i + \alpha(\theta)PRIVATE_i + \gamma(\theta)YEAR_i + X_i\psi(\theta) + \varepsilon_i(\theta)$$

The parameter $\beta^{DD}(\theta)$ evaluates the treatment effect of the pension reform on the household saving rate at the θth quantile. Compared with linear difference-in-differences estimation, identifying the quantile treatment effect by using quantile difference-in-differences requires a more stringent common trend assumption. It requires a common trend to exist in each quantile of the household saving rate distribution.

The results in Table 4.5 imply that the mandatory workplace pensions have had a significantly negative impact on the households at the bottom and the median of the saving rate distribution (10th percentile to 60th percentile) but little impact on the top saving rate quantile (above the 70th percentile). That is, the estimated reform impact shown in my main result is concentrated on those households with low to median saving rates.

This result may reveal that job sorting for workplace pensions may have existed among employees before the reform. That is, employees with a stronger preference for saving (e.g. people who would like to consume more in retirement) might choose jobs offering more generous pension plans. Before the reform, such workers may already have had workplace pensions but also high voluntary savings. In other words, these employees with employer pension contributions may have stayed in a relatively high quantile of the savings distribution. Hence, the expansion of workplace pensions brought about by the reform should have had less impact on these workers. This gives a possible explanation for the differential results seen in the top and bottom quantiles in Table 4.5. In addition, those with low household saving rate are likely to be liquidity constrained. Therefore, when employers help them contribute 6% of the wage to pension accounts, they are more likely to reduce their saving to smooth out their consumption, which provides another explanation for my findings.

4.8 Conclusion

This paper exploits the recent workplace pension reform in Taiwan as a natural experiment through which to investigate the impact of workplace pension provision on households' voluntary savings. My results suggest that this reform significantly reduces household saving rates by 2.06–2.45 percentage points on average. This implies that the average elasticity of substitution between workplace pension and households' voluntary savings is between -0.50 and -0.60. Moreover, to examine the reform's impact on the entire household saving rate distribution, I conducted unconditional quantile difference-in-differences estimation and found that most of the average policy effect is indeed concentrated in the bottom and median quantiles. This finding can be explained by employees' job-sorting behavior, with the savings-oriented attracted to workplace pensions before the reform. In general, I find that workplace pensions crowd out only half of household saving, which is similar to previous studies on workplace pensions (Engelhardt and Kumar, 2011) or public pension (Attanasio and Brugiavini, 2003; Attanasio and Rohwedder, 2003). Therefore, my results suggest that mandatory workplace pensions could be an effective policy instrument for raising employees' retirement wealth.

However, one important caveat should be noted in the interpretation of my results. The TSFIE data lack information about the pension coverage rate before the reform and the choice of pension scheme after the reform. Hence I know only the eligible group of employees and not the truly affected population. For this reason, my difference-in-differences estimates identify the intention-to-treat effect but not the average treatment effect on the treated. From the aggregate data, I find that the reform may have made 80% of private-sector employees newly eligible for workplace pensions when retiring, and that 84% of employees are covered by the new pension scheme after the reform. Hence, this pension reform actually affected the majority of private-sector employees, which may substantially mitigate the bias this data problem could have had on my estimates of the reform's impact. Nevertheless, it would still be worth linking administrative data from the government pension authority with TSFIE data to obtain more precise estimates of the pension-saving offset.

.

4.9 Figures

Figure 4.1: Workplace Pension Coverage in Private Sector: 2002–2007

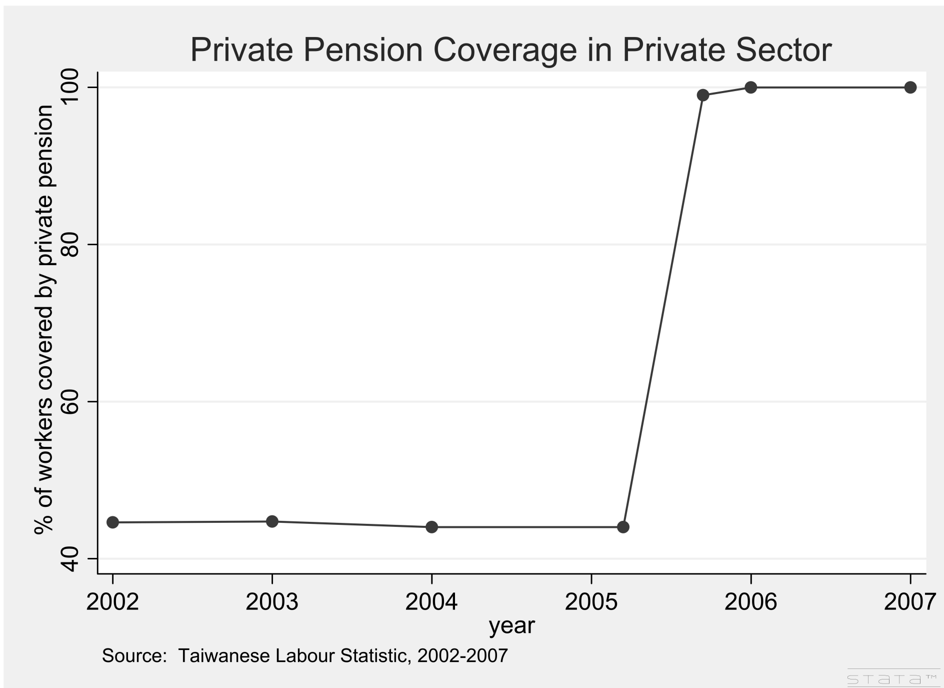

4.10 Tables

Table 4.1: Comparison Between New and Old Pension Systems

	New pension system	Old pension system
System	Defined contribution system	Defined benefit system
Law	Labor pension law	Labor standard law
Vesting period	Immediate vesting	The employees are required to stay in the same firm for 25 years or stay in the same firm for 15 years and become 55 years old
Employer's contribution	Mandatory rate: at least 6% of an employee's wage	Flexible rate: 2% to 15% of an employee's wage

Source: Taiwanese Council of Labor Affairs

Table 4.2: Descriptive Statistics

	Private sector			Public sector			
	Pre-reform (2002–2004)	Post-reform (2006–2008)	Diff	Pre-reform (2002–2004)	Post-reform (2006–2008)	Diff	Diff-in-Diff estimates
Saving rate	0.22	0.206	-0.014***	0.309	0.312	0.002	-0.016**
	(0.22)	(0.211)	[0.003]	(0.232)	(0.215)	[0.007]	[0.007]
Saving	27.431	25.557	-1.8741***	47.258	48.786	1.528	-3.402**
	(43.211)	(39.235)	[0.487]	(50.207)	(48.47)	[1.546]	[1.417]
Wage income	67.343	64.897	-2.446***	82.993	84.2	1.206	-3.652***
	(40.641)	(40.572)	[0.479]	(42.22)	(45.737)	[1.370]	[1.371]
Non-wage income	34.828	35.204	0.376	49.54	50.497	0.957	-0.581
(after-tax)	(44.236)	(38.609)	[0.490]	(43.425)	(40.351)	[1.316]	[1.391]
Consumption	74.739	74.544	-0.195	85.276	85.912	0.635	-0.831
	(37.31)	(35.534)	[0.430]	(36.852)	(35.125)	[1.129]	[1.217]
Head's age	37.998	38.477	0.478***	40.393	40.929	0.536**	-0.057
	(7.257)	(7.33)	[0.086]	(6.672)	(6.743)	[0.210]	[0.241]
Head's education	12.045	12.556	0.510***	14.267	14.646	0.379***	0.130
	(3.041)	(2.878)	[0.035]	(2.404)	(2.402)	[0.075]	[0.097]
Spouse's education	7.25	7.048	-0.202***	9.494	9.302	-0.191	-0.010
	(6.001)	(6.318)	[0.073]	(6.335)	(6.761)	[0.204]	[0.208]
male head	0.775	0.749	-0.025***	0.763	0.724	-0.039***	0.013
	(0.418)	(0.433)	[0.005]	(0.425)	(0.447)	[0.014]	[0.014]
# of above 18	2.586	2.572	-0.0140	2.304	2.286	-0.017	0.003
	(1.1)	(1.07)	[0.013]	(0.876)	(0.824)	[0.027]	[0.035]
# of below 18	1.224	1.101	-0.122***	1.364	1.247	-0.117***	-0.006
	(1.099)	(1.058)	[0.013]	(1.046)	(1.027)	[0.032]	[0.036]
# of above 65	0.27	0.311	0.041***	0.229	0.241	0.012	0.029
	(0.558)	(0.597)	[0.007]	(0.529)	(0.546)	[0.017]	[0.019]
# of working	1.592	1.587	-0.005	1.352	1.347	-0.006	0.001
	(0.723)	(0.706)	[0.008]	(0.497)	(0.487)	[0.015]	[0.023]
Southern Taiwan	0.178	0.211	0.033***	0.184	0.20	0.016	0.016
	(0.383)	[0.408]	[0.005]	[0.388]	[0.40]	[0.012]	[0.013]
Middle Taiwan	0.183	0.207	0.024***	0.213	0.205	-0.008	0.032**
	(0.387)	(0.405)	[0.005]	(0.409)	(0.404)	[0.013]	[0.013]
Northern Taiwan	0.308	0.244	-0.063***	0.219	0.20	-0.019	-0.044***
	(0.205)	(0.430)	[0.005]	(0.414)	(0.40)	[0.013]	[0.015]
Observations	14,304	14,425	28,729	2,299	1,841	4,140	32,869

Note: Saving, wage income, non-wage income, and consumption are scaled in thousands of 2007 New Taiwan Dollars. Household disposable income is the sum of wage income and non-wage income(after-tax). The 2007 exchange rate is 1 US Dollars = 32.4 New Taiwan Dollars. Southern Taiwan, Middle Taiwan, and Northern Taiwan are the living area dummies. Standard errors in brackets, *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 4.2: Descriptive Statistics

	Private sector			Public sector			
	Pre-reform (2002–2004)	Post-reform (2006–2008)	Diff	Pre-reform (2002–2004)	Post-reform (2006–2008)	Diff	Diff-in-Diff estimates
Agriculture	0.015	0.014	-0.001	0.006	0.005	-0.0012	0.001
	(0.12)	(0.116)	[0.001]	(0.078)	(0.07)	[0.002]	[0.004]
Manufacturing	0.515	0.506	-0.008	0.089	0.084	-0.0051	-0.003
	(0.5)	(0.5)	[0.006]	(0.284)	(0.277)	[0.009]	[0.016]
Service	0.471	0.48	0.009	0.905	0.911	0.0063	0.003
	(0.499)	(0.5)	[0.006]	(0.293)	(0.284)	[0.009]	[0.016]
Profession	0.383	0.396	0.013**	0.52	0.526	0.0056	0.007
	(0.486)	(0.489)	[0.006]	(0.5)	(0.499)	[0.016]	[0.016]
White collar	0.357	0.35	-0.007	0.305	0.311	0.0059	-0.013
	(0.479)	(0.477)	[0.006]	(0.461)	(0.463)	[0.014]	[0.016]
Blue collar	0.26	0.254	-0.006	0.174	0.163	-0.0115	0 .005
	(0.439)	(0.435)	[0.005]	(0.38)	(0.369)	[0.012]	[0.014]
Own house	0.84	0.857	0.018***	0.89	0.902	0.0123	0.006
	(0.367)	(0.35)	[0.004]	(0.313)	(0.297)	[0.010]	[0.012]
Housing size	40.484	42.314	1.830***	44.776	45.705	0.9290	0.901
	(20.253)	(21.743)	[0.248]	(21.577)	(22.229)	[0.684]	[0.706]
Observations	14,304	14,425	28,729	2,299	1,841	4,140	32,869

Note: household size is meausred in square footage. 10 square feet is equal to 3.3057 square meters. Standard errors in brackets, *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 4.3: The Effect of Mandatory Private Pension on Household Voluntory Saving

	Saving rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Period	2002–2008	2002–2008	2002–2008	2002–2008	2000–2004	2002–2008	2002–2008
Cohort	20–50	20–50	20–50	20–50	20–50	20–50	51–65
Treament group	private	private	private	private	private	banking	private
Pension effect $\widehat{\beta}^{DD}$	-0.0162**	-0.0202***	-0.0201***	-0.0206***	-0.0053	-0.0050	0.0066
	[0.006]	[0.005]	[0.006]	[0.006]	[0.003]	[0.008]	[0.008]
Baseline mean	0.22	0.22	0.22	0.22	0.26	0.20	0.29
Equal to (4)	accept	accept	accept		reject	reject	reject
Family chacteristic		$\surd$	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$
Industry & occupation			$\surd$	$\surd$	$\surd$	$\surd$	$\surd$
Household wealth				$\surd$	$\surd$	$\surd$	$\surd$
observation	32,869	32,869	32,869	32,869	22,422	5,924	8,593
R^2	0.023	0.225	0.230	0.250	0.235	0.277	0.298

Note: Family characteristic: head's age, age square, education, gender; spouse's education; # of children under 18, # of members over 65, # of members above 18, # of working members, and living county dummies Industry & occupation : head's industry and occupation Household wealth: household total non-wage income; dummy for indicating having their own house; housing size. Baseline mean is the saving rate for private sector households during pre-reform period (2002–2004). Standard errors clustered on sector/year in brackets, *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 4.4: Empirical Specification Checks

	Saving rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Standard error sector/period cluster	Statistical Inference block bootstrap	Statistical Inference wild bootstrap	Dependent var. ln(Y)-ln(C)	Treatment group head's sectoral choice	Different sample period	Control earned income
$\widehat{\beta}^{DD}$	-0.0206***	-0.0206***	-0.0206***	-0.0250**	-0.0137**	-0.0214***	-0.0138**
	[0.001]			[0.010]	[0.005]	[0.006]	[0.006]
p-value	(0.000)	(0.001)	(0.007)	(0.029)	(0.015)	(0.009)	(0.041)
observation	32,869	32,869	32,869	32,869	34,707	27,302	32,869
R^2	0.250			0.260	0.255	0.254	0.255

Note: Family chacteristic: head's age, age square, education, gender; spouse's education; # of children under 18, # of members over 65, # of members above 18, # of working members, and living county dummies Industry & occupation: head's industry and occupation Household wealth: household total non-wage income; dummy for indicating having their own house; housing size. Standard errors in block bootstrap and wild bootstrap are calculated by using 999 random repetitions. Standard errors in brackets, *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table 4.5: Quantile DD Results

	Quantile of saving rate								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Quantile	10th	20th	30th	40th	50th	60th	70th	80th	90th
Pension effect $\widehat{\beta}^{QDD}$	-0.0394***	-0.0289***	-0.0254***	-0.0242***	-0.0212***	-0.0165**	-0.0106	-0.0018	-0.0040
	[0.011]	[0.007]	[0.007]	[0.007]	[0.008]	[0.008]	[0.009]	[0.010]	[0.013]
observation	32,869	32,869	32,869	32,869	32,869	32,869	32,869	32,869	32,869
R^2	0.080	0.128	0.162	0.181	0.192	0.191	0.176	0.154	0.106

Note: Family characteristic: head's age, age square, education, gender; spouse's education; # of children under 18, # of members over 65, # of members above 18, # of working members, and living county dummies Industry & occupation : head's industry and occupation Household wealth: household total non-wage income; dummy for indicating having their own house; housing size. Block standard errors in brackets (cluster on sector/year), which are calculated by using 999 random repetitions. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Bibliography

Adams, N. (2014). Retirement confidence survey. Employee Benefit Research Institute.

Adams, W., L. Einav, and J. Levin (2009). Liquidity constraints and imperfect information in subprime lending. *American Economic Review 99*(1), 49–84.

Aguila, E. (2011). Personal retirement accounts and saving. *American Economic Journal: Economic Policy 3*(4), 1–21.

Alessie, R., A. Kapteyn, and F. Klijn (1997). Mandatory pensions and personal savings in the Netherlands. *De Economist 145*(3), 291–324.

Almond, D. (2006). Is the 1918 influenza pandemic over? Long-term effects of in utero influenza exposure in the post-1940 U.S. population. *Journal of Political Economy 114*(4), 672–712.

Almond, D., J. Doyle, A. Kowalski, and H. Williams (2011). Estimating marginal returns to medical care: Evidence from at-risk newborns. *Quarterly Journal of Economics 125*(2), 591–634.

Anderson, M., C. Dobkin, and T. Gross (2012). The effect of health insurance coverage on the use of medical services. *American Economic Journal: Economic Policy 4*(1), 1–27.

Andrietti, V. and V. A. Hildebrand (2006). Evaluating pension portability reforms: The Tax Reform Act of 1986 as a natural experiment. *Working paper*.

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001). The hyberbolic consumption model: Calibration, simulation, and empirical evaluation. *Journal of Economic Perspectives 15*(3), 47–68.

Angrist, J. D. and J.-S. Pischke (2009). Mostly harmless econometrics: An empiricist's companion. Princeton University Press.

Aron-Dine, A., L. Einav, , and A. Finkelstein (2013). The RAND health insurance experiment, three decades later. *Journal of Economic Perspectives 27*(1), 197—222.

Attanasio, O. P. and A. Brugiavini (2003). Social security and households' saving. *Quarterly Journal of Economics 118*(3), 1075–1119.

Attanasio, O. P. and S. Rohwedder (2003). Pension wealth and household saving: Evidence from pension reforms in the United Kingdom. *American Economic Review 93*(5), 1499–1521.

Baicker, K., S. Mullainathan, and J. Schwartzstein (2013). Behavioral hazard in health insurance. *Harvard Working Paper*.

Baker, M., J. Gruber, and K. Milligan (2008, August). Universal child care, maternal labor supply, and family well-being. *Journal of Political Economy 116*(4), 709–745.

Banerjee, A. V. and S. Mullainathan (2010). The shape of temptation: Implications for the economic lives of the poor. NBER Working Paper No. 15973.

Banerjee, S. (2011). Does social security affect household saving? *Working paper*.

Barreca, A. I., M. Guldi, J. M. Lindo, and G. R. Waddell (2011). Robust nonparametric confidence intervals for Regression-Discontinuity designs. *The Quarterly Journal of Economics 126*(4), 2117–2123.

Barrow, L. and L. McGranahan (2001). The effects of the Earned Income Tax Credit on the seasonality of household expenditures. *National Tax Journal 53*(4), 1211–1244.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

Berube, A., A. Kim, B. Forman, and M. Burns (2002). The price of paying taxes: How tax preparation and refund loan fees erode the benefits of the EITC. Brookings Institution. Washington, DC, May.

Bharadwaj, P., K. V. Løken, and C. Neilson (2013). Early life health interventions and academic achievement. *American Economic Review 103*(5), 1862–1891.

Bitler, M., H. Hoynes, and E. Kuka (2014). Do in-work tax credits serve as a safety net? *NBER Working Paper No. 19785*.

Blau, F. D. and L. M. Kahn (2007). Changes in the labor supply behavior of married women: 1980-2000. *Journal of Labor Economics 25*(3), 393–438.

Bloom, H. S. (1984, April). Accounting for no-shows in experimental evaluation designs. *Evaluation Review 8*(2), 225–246.

Blundell, R. W., L. Pistaferri, and I. S. Eksten (2014). Consumption inequality and family labor supply. *Working Paper*.

Browning, M. and M. D. Collado (2001). The response of expenditures to anticipated income changes: Panel data estimates. *American Economic review 91*(3), 681–692.

Cagan, P. (1965). The effect of pension plans on aggregate saving: Evidence from a sample survey. *NBER Occasional Paper 95*.

Cameron, C., J. Gelbach, and D. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics 90*(3), 414—427.

Card, D., R. Chetty, and A. Weber (2007). Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. *Quarterly Journal of Economics 122*(4), 1511–1560.

Card, D., C. Dobkin, and N. Maestas (2008). The impact of nearly universal insurance coverage on health care utilization: Evidence from medicare. *American Economic Review 98*(5), 2242–2258.

Card, D., C. Dobkin, and N. Maestas (2009). Does medicare save lives? *The Quarterly Journal of Economics 124*(2), 597–636.

Case, A., A. Fertig, and C. Paxson (2005). The lasting impact of childhood health and circumstance. *Journal of Health Economics 24*(2), 365–389.

Cattaneo, M. D., S. Calonico, and R. Titiunik (2013). Robust nonparametric confidence intervals for Regression-Discontinuity designs. *Working paper*.

Center of Disease and Control (2013). *Vaccination Schedule in Taiwan*. Center of Disease and Control.

Chandra, A., J. Gruber, and R. McKnight (2010a). Patient cost-sharing and hospitalization offsets in the elderly. *American Economic Review 100*(1), 193–213.

Chandra, A., J. Gruber, and R. McKnight (2010b). Patient cost sharing in low income populations. *American Economic Review 100*(2), 303–308.

Chandra, A., J. Gruber, and R. McKnight (2014). The impact of patient cost-sharing on low-income populations: Evidence from Massachusetts. *Journal of Health Economics 33*(1), 57–66.

Cherkin, D., L. Grothaus, and E. Wagner (1989). The effect of office visit copayments on utilization in a health maintenance organization. *Medical Care 27*, 1036–1045.

Chetty, R. (2008). Moral hazard vs. liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark. *Quarterly Journal of Economics 129*(3), 1141–1219.

Chetty, R., J. N. Friedman, and E. Saez (2013a). Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *American Economic Review 103*(7), 2683–2721.

Chetty, R., J. N. Friedman, and E. Saez (2013b). Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *American Economic Review 103*(7), 2683–2721.

Chetty, R. and E. Saez (2013). Teaching the tax code: Earnings responses to an experiment with EITC recipients. *American Economic Journal: Applied Economics 5*(1), 1–31.

Chou, S. Y., J. T. Liu, and J. K. Hammitt (2003). National health insurance and precautionary saving: Evidence from Taiwan. *Journal of Public Economics 87*, 1873–1894.

Cullen, J. B. and J. Gruber (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics 18*(3), 546–572.

Currie, J. (2009). Healthy, wealthy, and wise: Socioeconomic status, poor health in childhood, and human capital development. *Journal of Economic Literature 47*(1), 87–122.

Currie, J. and B. Madrian (1999). Health, health insurance, and the labor market. *Handbook of Labor Economics 3*(2), 365–389.

Deaton, A. and C. Paxson (2000). Growth and saving among individuals and households. *The Review of Economics and Statistics 82*(2), 212–225.

Del Boca, D. and A. Lusardi (2003). Credit Market Constraints and labor market decisions. *Labour Economics 10*(1), 681–703.

Edmonds, E. V. (2006). Child labor and schooling responses to anticipated income in South Africa. *Journal of Development Economics 81*(2), 386–414.

Edwards, R. D. (2004). Macroeconomic implications of the Earned Income Tax Credit. *National Tax Journal 57*(1), 45–65.

Eissa, N. and H. Hoynes (2011). Redistribution and tax expenditures: The Earned Income Tax Credit. *National Tax Journal 64*(2), 689–730.

Eissa, N. and H. W. Hoynes (2004). Taxes and the labor market participation of married couples: The Earned Income Tax Credit. *Journal of Public Economics 88*(9), 1931–1958.

Eissa, N. and H. W. Hoynes (2006). Behavioral responses to taxes: Lessons from the EITC and labor supply. *Tax Policy and the Economy 20*, 74–110.

Eissa, N. and J. Liebman (1996). Labor supply response to the Earned Income Tax Credit. *Quarterly Journal of Economics 111*(2), 605–637.

Engelhardt, G. and A. Kumar (2011). Pensions and household wealth accumulation. *Journal of Human Resources 46*(1), 203—236.

Engelhardt, G. V. (2001). Have 401(k)s raised household saving? Evidence from the health and retirement study.

Euwals, R. (2000). Do mandatory pensions decrease household savings? Evidence for the Netherlands. *De Economist 148*(5), 643–670.

Feldstein, M. S. (1974). Social security, induced retirement, and aggregate capital accumulation. *Journal of Political Economy 85*(5), 905—926.

Feng, J., L. He, and H. Sato (2011). Public pension and household saving: Evidence from urban China. *Journal of Comparative Economics 39*(4), 470–485.

Fernandez, F. and V. Saldarriaga (2014). Do benefit recipients change their labor supply after receiving the cash transfer? Evidence from the Peruvian Juntos program. *IZA Journal of Labor & Development 3*(2), 1–30.

Firpo, S., N. M. Fortin, and T. Lemieux (2009). Unconditional quantile regressions. *Econometrica 77*(3), 953–973.

Fortin, N. M. (1995). Allocation inexibilities, female labor supply, and housing assets accumulations: Are women working to pay the mortgage? *Journal of Labor Economics 13*(1), 524–557.

Gale, W., J. Shoven, and M. Warshawsky (2005). The evolving pension system: Trends, effects, and proposals for reform. *Brookings Institution.*

Gale, W. G. (1998). The effects of pension on household wealth: A reevaluation of theory and evidence. *Journal of Political Economy 82*(5), 905–926.

Gale, W. G. and J. K. Scholz (1994). IRAs and household saving. *American Economic Review 84*(5), 1233–1260.

GAO (1992). Earned Income Tax Credit: Advance payment option is not widely known or understood by the public. *Report to Congressional Committees.*

Gelber, A. M. and J. W. Mitchell (2012). Taxes and time allocation: Evidence from single women and men. *Review of Economic Studies 79*(3), 863–897.

Goux, D., E. M. and B. Petrongolo (2014). Worktime regulations and spousal labor supply. *American Economic Review 104*(1), 252–276.

Grogger, J. (2003). The effects of time limits, the EITC, and other policy changes on welfare use, work, and income among female-headed families. *Review of Economics and Statistics 85*(2), 394–408.

Gustman, A. L. and T. L. Steinmeier (1999). Effects of pensions on savings: Analysis with data from the health and retirement study. *Carnegie-Rochester Series on Public Policy 50*(1), 271—326.

Heckman, J. J. and T. E. MaCurdy (1980). A life cycle model of female labour supply. *Review of Economic Studies 47*(1), 47–74.

Heim, B. T. (2007). The incredible shrinking elasticities: Married female labor supply, 1978-2002. *The Journal of Human Resources 42*(4), 881–918.

Hemming, R. and R. Harvey (1983). Occupational pension scheme membership and retirement saving. *Economic Journal 93*, 128–144.

Holt, S. (2009). Beyond lump sum: Periodic payment of the Earned Income Tax Credit. *Working Paper*.

Hotz, V. J. and J. K. Scholz (2003). The Earned Income Tax Credit. *Means-Tested Transfer Programs in the United States*. Edited by Robert Moffitt. Chicago: University of Chicago Press.

Hotz, V. J. and J. K. Scholz (2006). Examining the effect of the Earned Income Tax Credit on the labor market participation of families on welfare. *National Bureau of Economic Research Working Paper 11968*.

Hoynes, H. W., D. L. Miller, and D. Simon (2014). Income, the Earned Income Tax Credit, and infant health. *American Economic Journal: Economic Policy*, forthcoming.

Hsieh, C.-T. (2003). Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund. *American Economic review 93*(1), 397–405.

Hubbard, R. G. (1986). Pension wealth and individual saving. *Journal of Money, Credit, and Banking 18*(2), 167–178.

Imbens, G. and K. Kalyanaraman (2012). Optimal bandwidth choice for the Regression Discontinuity estimator. *The Review of Economic Studies 79*(3), 933–959.

Imbens, G. W. and J. D. Angrist (1994, Mar). Identification and estimation of local average treatment effects. *Econometrica 62*(2), 467–475.

Jacob, B. A. and J. Ludwig (2012). The effects of housing assistance on labor supply: Evidence from a voucher lottery. *American Economic Review 102*(1), 272–304.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics 2*, 479–506.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic review 96*(5), 1589—1610.

Johnson, D. S., J. A. Parker, and N. S. Souleles (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic review 103*(6), 2530—2553.

Jones, D. (2010). Information, preferences, and public benefit participation: Experimental evidence from the advance EITC and 401(k) savings. *American Economic Journal: Applied Economics 2*(2), 147–163.

Jones, D. (2012). Inertia and overwithholding: Explaining the prevalence of income tax refunds. *American Economic Journal: Economic Policy 4*(1), 158–185.

Katona, G. (1965). Private pensions and individual savings. *Survey Research Center, Institute for Social Research. 40.*

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33—50.

Kohara, M. (2010). The response of Japanese wives labor supply to husbands job loss. *Journal of Population Economics 23*(4), 1133–1149.

Kotlikoff, L. J. (1979, June). Testing the theory of social security and life cycle accumulation. *American Economic Review 69*(3), 396–410.

LaLumia, S. (2013). The EITC, tax refunds, and unemployment spells. *American Economic Journal: Economic Policy 5*(2), 188–221.

Lee, D. S. and T. Lemieux (2010). Regression Discontinuity designs in economics. *Journal of Economic Literature 48*(2), 281–355.

Leibowitz, Manning, Keeler, Duan, Lohr, and Newhouse (1985). Effect of cost-sharing on the use of medical services by children: Interim results from a randomized controlled trial. *Pediatrics 75*(5), 942–951.

Lemieux, T. and K. Milligan (2008). Incentive effects of social assistance: A Regression Discontinuity approach. *Journal of Econometrics 142*(2), 807–828.

Lien, H.-M., S.-Y. Chou, and J.-T. Liu (2008). Hospital ownership and performance: Evidence from stroke and cardiac treatment in Taiwan. *Journal of Health Economics 27*(5), 1208–1223.

Looney, A. and M. Singhal (2006). The effect of anticipated tax changes on intertemporal labor supply and the realization of taxable income. *NBER Working Paper No. 12417*.

Ludwig, J. and D. Miller (2007). Does head start improve children’ s life chances? Evidence from a Regression Discontinuity design. *Quarterly Journal of Economics 122*(2), 159–208.

Lundberg, S. (1980). The added worker effect. *Journal of Labor Economics 3*(1), 11–37.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz (1981). Some interim results from a controlled trial of cost sharing in health insurance. *New England Journal of Medicine 305*(1), 1501—1507.

McClelland, R. and S. Mok (2012). A review of recent research on labor supply elasticities. *Congressional Budget Office Working Paper*.

McGranahan, L. and D. W. Schanzenbach (2014). The Earned Income Tax Credit and food consumption patterns. *Working Paper*.

Meyer, B. D. (1995). Natural and quasi-experiments in economics. *Journal of Business and Economic Statistics 13*, 151–161.

Meyer, B. D. (2010). The effects of the Earned Income Tax Credit and recent reforms. *Tax Policy and the Economy 24*, 153–180. Edited by Jeffrey R. Brown. Cambridge, MA: MIT Press.

Meyer, B. D. and D. T. Rosenbaum (2001). Welfare, the Earned Income Tax Credit, and the labor supply of single mothers. *Quarterly Journal of Economics 116*(3), 1063–1114.

Mian, A., K. Rao, and A. Sufi (2014). Household balance sheets, consumption, and the economics slump. *Quarterly Journal of Economics*, forthcoming.

Mian, A. and A. Sufi (2011). House prices, home equity based borrowing, and the U.S. household leverage crisis. *American Economic Review 101*, 2132–2156.

Mincer, J. (1962). Labor force participation of married women: A study of labor supply. *Aspects of Labor Economics*, 63–97. Edited by H. G. Lewis. Princeton N.J.: National Bureau of Economic Research, Princeton University Press.

Mullainathan, S., M. Bertrand, and E. Duflo (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics 119*(1), 249–275.

Munnell, A. H. (1976). Private pensions and savings: New evidence. *Journal of Political Economy 84*(5), 1013–1032.

National Health Insurance Administration (2012). *National Health Insurance Research Database codebook*. National Health Insurance Administration.

Niedzwiecki, M. J. (2013). (Re)funding health care spending: The timing of EITC refunds, liquidity, and investment in health.

OECD (2005). Pensions Glossary. *Organisation for Economic Co-operation and Development*.

OECD (2012). OECD Pensions Outlook 2012. *Organisation for Economic Co-operation and Development*.

Parker, J. A. (1999). The reaction of household consumption to predictable changes in social security taxes. *American Economic review 89*(4), 959–973.

Parker, J. A. (2014). Why don't households smooth consumption? Evidence from a 25 million dollar experiment. *Manuscript*.

Paxson, C. H. (1993). Consumption and income seasonality in Thailand. *Journal of Political Economy 101*(1), 39–72.

Pence, K. (2001). 401(k)s and household saving: New evidence from the survey of consumer finances.

Poterba, J. M., S. F. Venti, and D. A. Wise (1995). Do 401(k) contributions crowd out other personal saving? *Journal of Public Economics 58*(1), 1–32.

Rice, T. and K. Y. Matsuoka (2004). The impact of cost-sharing on appropriate utilization and health status: A review of the literature on seniors. *Medical Care Research and Review 61*(4), 415—452.

Romich, J. L. and T. Weisner (2000). How families view and use the EITC: Advance payment versus lump sum delivery. *National Tax Journal 53*(4), 1245–1266.

Rothstein, J. (2010). Is the EITC as good as an NIT? Conditional cash transfers and tax incidence. *American Economic Journal: Economic Policy 2*(1), 177–208.

Saez, E. (2003). The effect of marginal tax rates on income: A panel study of bracket creep. *Journal of Public Economics 87*(5), 1231–1258.

Selby, J., B. Fireman, and B. Swain. (1996). Effect of a co-payment on use of the emergency department in a health maintenance organization. *New England Journal of Medicine 334*(1), 635—641.

Selden, T. M., G. M. Kenney, M. S. Pantell, and J. Ruhter (2009). Cost sharing in Medicaid and CHIP: How does it affect out-of-pocket spending? *Health Affairs 28*(4), 607–619.

Sen, Blackburn, Morrisey, Kilgore, Becker, Caldwell, and Menachemi (2012). Did copayment changes reduce health service utilization among CHIP enrollees? Evidence from Alabama. *Health Services Research 47*(4), 1603—1620.

Shigeoka, H. (2014). The effect of patient cost-sharing on utilization, health and risk protection. *American Economic Review*, forthcoming.

Sisson, J. and K. Short (2001). Measuring and modeling taxes in the survey of income and program participation. *Census Bureau, Working Paper*.

Smeeding, T. M., K. R. Phillips, and M. O' Connor (2000). The EITC: Expectation, knowledge, use, and economic and social mobility. *National Tax Journal 53*(4), 1187–1210.

Souleles, N. S. (1999). The response of household consumption to income tax refunds. *American Economic review 89*(4), 947–958.

Stephens, Melvin, J. (2003). 3rd of tha month' : Do social security recipients smooth consumption between checks? *American Economic review 93*(1), 406–422.

Stephens, Melvin, J. (2006). Paycheque receipt and the timing of consumption. *Economic Journal 116*(513), 680–701.

Stephens, M. and T. Unayama (2011). The consumption response to seasonal income: Evidence from Japanese public pension benefits. *American Economic Journal: Applied Economics 3*(4), 86–118.

Summers, L. H. (1989). Some simple economics of mandated benefits. *American Economic Review 79*, 177–183.

Taiwanese Council of Labor Affairs (2002-2007). Taiwanese Labour Statistic. *Year Book*.

Tapia, W. (2008). Description of private pension systems. *OECD Working Papers on Insurance and Private Pensions*.

Wu, C. C. (2012). The party's over for quickie tax loans: But traps remain for unwary taxpayers. *Working Paper*.

Yang, T.-T. and M.-C. Luoh (2009). Who pays pension? The impact of new pension scheme on labour wage. *Academia Economic Papers 37*(3), 339 – 368.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy 97*(2), 305–346.

Appendix A

Appendix to Chapter 2

A.1 Appendix Figures

Figure A.1: The Impact of EITC on Intra-Year Labor Supply Pattern (including Pre-Trend): Married Women

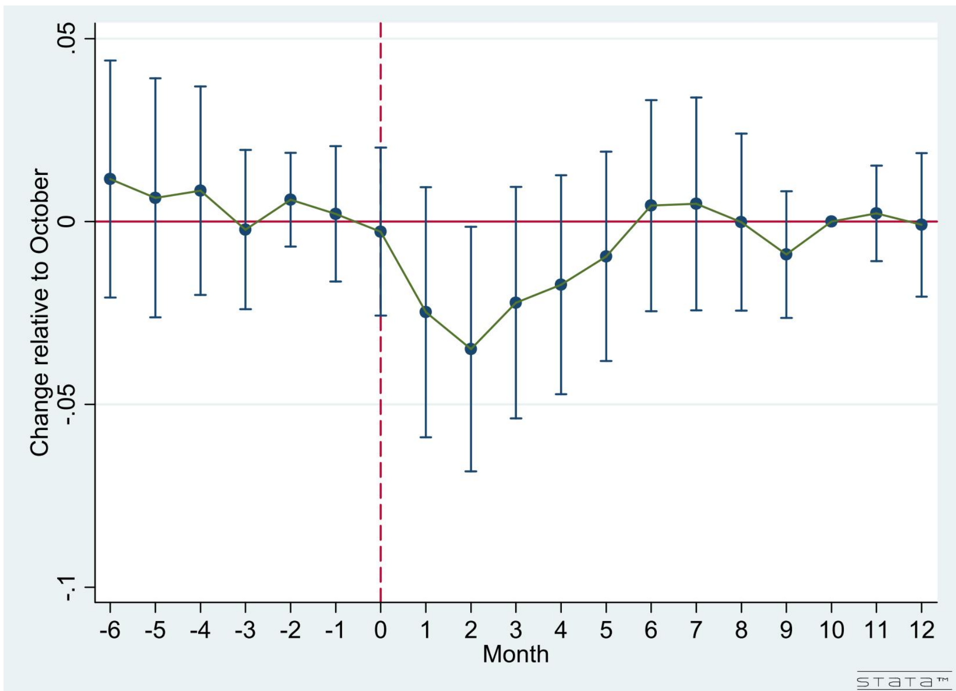

Notes: This figure shows coefficients on $EITC_{it} \times M$ (M includes twelve months plus last seven months in the previous year) and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to married women. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than $36,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than $36,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above $36,000 and below $40,000. (3) childless individuals that have incomes greater than $36,000 and below $40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below $36,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than $36,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Figure A.2: The Impact of EITC on Intra-Year Labor Supply Pattern (including Pre-Trend): Married Men

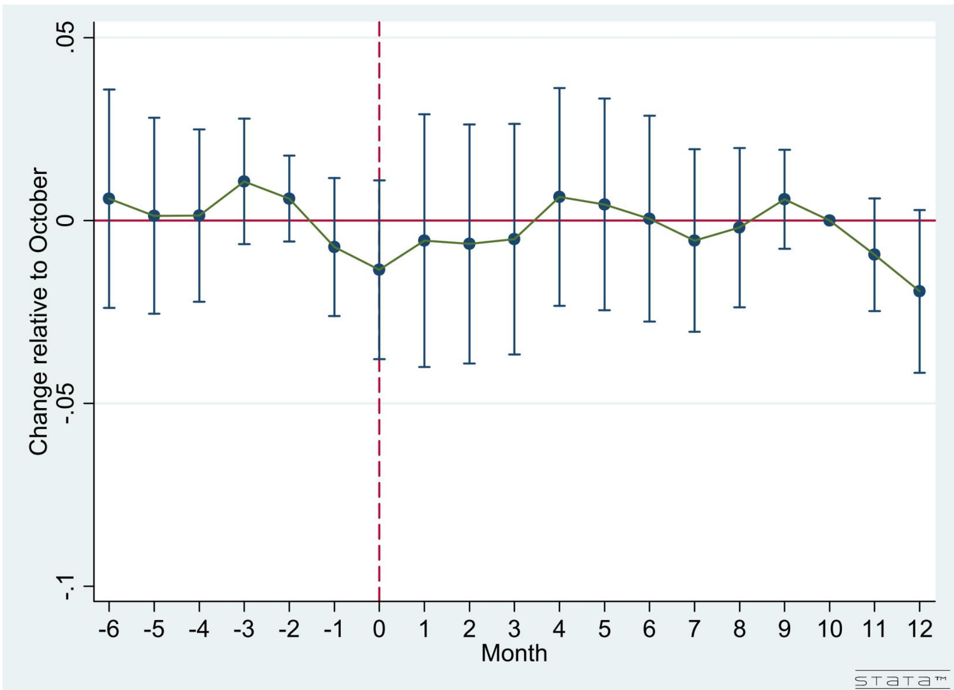

Notes: This figure shows coefficients on $EITC_{it} \times M$ (M includes twelve months plus last seven months in the previous year) and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to married men. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than \$36,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than \$36,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above \$36,000 and below \$40,000. (3) childless individuals that have incomes greater than \$36,000 and below \$40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below \$36,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than \$36,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Figure A.3: The Impact of EITC on Intra-Year Labor Supply Pattern (including Pre-Trend): Single Women

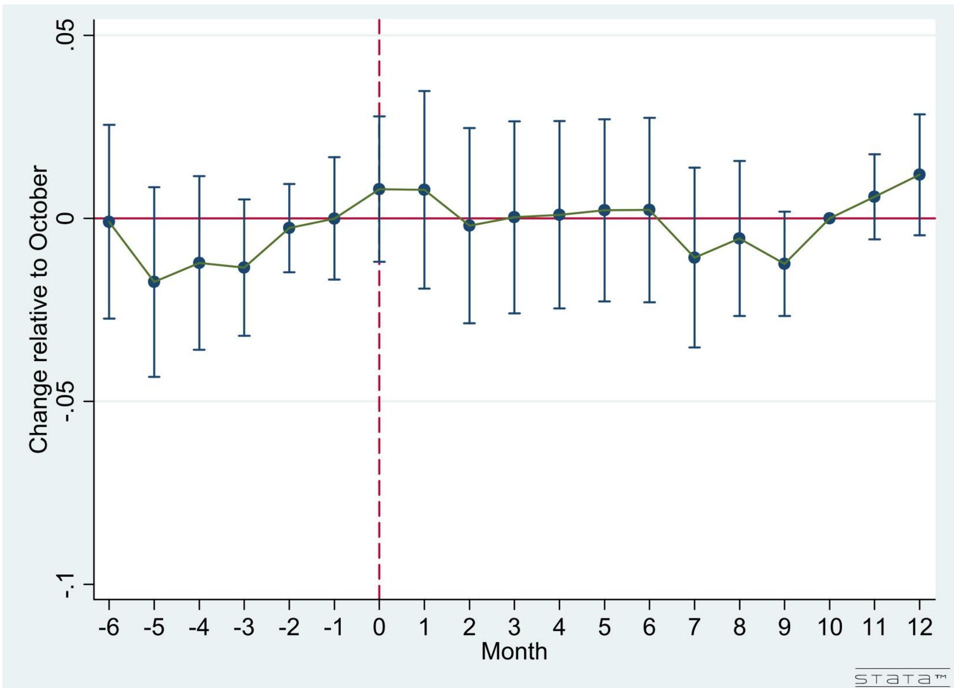

Notes: This figure shows coefficients on $EITC_{it} \times M$ (M includes twelve months plus last seven months in the previous year) and associated 95% confidence interval from specification 2.2 where the dependent variable L is the share of weeks worked in a month defined as number of working weeks divided by total number of weeks in a month. Therefore, $L = 1$ if working for the full month, $L = 0$ if not working for the full month, and $0 < L < 1$ if working for partial month. The estimated sample is restricted to single women. The dependent variable is regressed on the interaction terms between indicator for treatment group $EITC$ and 11 month dummies (October is the omitted month) M. The treatment group consists of those individuals that have one or more qualifying children and family income during tax year greater than zero and less than \$33,000. The comparison groups comprise (1) those individuals that have family income during tax year greater than zero and less than \$33,000 but have no qualifying child. (2) those individuals with one or more qualifying children but whose annual income is just above \$33,000 and below \$40,000. (3) childless individuals that have incomes greater than \$33,000 and below \$40,000. All dollar values are measured in 2007 dollars. The regression controls for treatment group dummy, an indicator for individuals with one or more qualifying children, an indicator for individuals with family income greater zero and below \$33,000, month fixed effect for those who have qualifying children, month fixed effect for those who have family income during tax year less than \$33,000, month fixed effect, individual fixed effect, year fixed effect, state fixed effect, monthly state unemployment rate, state specific time trend (quadratic), an indicator for interviewing month, educational attainment, number of children under 18, age, industry fixed effect, industry specific time trend (quadratic), family wealth, a dummy indicating that the individual worked part-time in the previous year, and month fixed effect specific to part-time workers.

Appendix B

Appendix to Chapter 3

B.1 Appendix Tables

Table B.1: Placebo Test for Other Age Cutoff

Panel A: Log(outpatient expenditure)			
Cutoff Age (days)	Coefficient on cutoff	Cutoff Age (days)	Coefficient on cutoff
886	0.66	1186	-0.63
	[0.42]		[0.39]
916	0.09	1216	-0.31
	[0.37]		[0.42]
946	-0.55	1246	0.85*
	[0.39]		[0.50]
976	-0.46	1276	-0.59
	[0.38]		[0.42]
1006	0.01	1306	-0.22
	[0.38]		[0.42]
1096	-6.90***	1336	0.51
(or 1095)	[0.49]		[0.44]
Panel B: Log(outpatient visits)			
Cutoff Age (days)	Coefficient on cutoff	Cutoff Age (days)	Coefficient on cutoff
886	0.24	1186	-0.80***
	[0.25]		[0.30]
916	-0.21	1216	-0.23
	[0.29]		[0.27]
946	-0.21	1246	0.59*
	[0.27]		[0.30]
976	-0.26	1276	-0.60**
	[0.25]		[0.26]
1006	-0.26	1306	-0.12
	[0.22]		[0.31]
1096	-4.73***	1336	0.19
(or 1095)	[0.31]		[0.31]

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure and the log of the total number of outpatient visits at each day of age. Column (1) and (3) indicates different cutoff age (measured in days) used in RD estimation. Note that 1096th (or 1095th) age day is the 3rd birthday and its estimate is corresponding to our main result in Table 3.6. Column (2) and (4) present estimated regression discontinuities of each interested outcome using data within 90 days before and after the 3rd birthday and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). All coefficients on $Age3$ and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table B.2: Sensitivity to Bandwidth and Polynomial Selection in Parametric RD Regressions

	Log(outpatient expenditure)					
Bandwidth (days)	60	120	180	240	300	360
Polynominal						
1	-6.69***	-6.19***	-5.54***	-5.10***	-4.54***	-4.65***
	[0.48]	[0.33]	[0.28]	[0.24]	[0.23]	[0.20]
2	-6.58***	-6.90***	-6.61***	-6.24***	-6.06***	-5.29***
	[0.74]	[0.51]	[0.40]	[0.37]	[0.32]	[0.30]
3	-7.07***	-6.68***	-7.04***	-6.98***	-6.85***	-6.94***
	[1.11]	[0.70]	[0.56]	[0.47]	[0.42]	[0.40]
	Log(outpatient visits)					
Bandwidth (days)	60	120	180	240	300	360
Polynominal						
1	-4.55***	-3.92***	-3.39***	-2.88***	-2.35***	-2.52***
	[0.34]	[0.24]	[0.20]	[0.18]	[0.17]	[0.15]
2	-4.33***	-4.97***	-4.36***	-4.12***	-3.89***	-3.04***
	[0.53]	[0.37]	[0.29]	[0.26]	[0.23]	[0.23]
3	-4.86***	-4.41***	-5.07***	-4.72***	-4.68***	-4.84***
	[0.83]	[0.49]	[0.41]	[0.33]	[0.30]	[0.29]

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure and the log of the total number of outpatient visits at each day of age. Each row indicates different order of polynomials used in RD estimation and each column denotes various bandwidth choice. We obtain RD estimates using OLS regression with uniform kernel function (similar to the parametric estimation in Table 3.6). Robust standard error in parentheses. All coefficients on $Age3$ and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table B.3: Sensitivity to Bandwidth Selector and Kernel Function Selection in Nonparametric RD Regressions

	Log(outpatient expenditure)			Log(outpatient visits)		
Bandwidth selector	CCT	IK	CV	CCT	IK	CV
Kernel function						
Triangular	-6.64***	-6.63***	-6.56***	-4.48***	-4.51***	-4.45***
	[0.48]	[0.44]	[0.40]	[0.39]	[0.35]	[0.45]
Bandwidth	81	89	105	67	79	54
Uniform	-6.68***	-6.69***	-6.58***	-4.46***	-4.46***	-4.40***
	[0.47]	[0.46]	[0.52]	[0.36]	[0.36]	[0.37]
Bandwidth	65	66	54	56	56	54
Epanechnikov	-6.64***	-6.64***	-6.64***	-4.45***	-4.49***	-4.43***
	[0.47]	[0.44]	[0.42]	[0.39]	[0.35]	[0.42]
Bandwidth	75	82	88	61	70	54

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure and the log of the total number of outpatient visits at each day of age. Each row indicates the specific kernel function used in nonparametric RD estimation and each column denotes the optimal bandwidth selector for choosing bandwidth. CCT is an optimal bandwidth selection method proposed by Matias D. Cattaneo, Sebastian Calonico and Rocio Titiunik (2013). IK is an optimal bandwidth selection procedure proposed by imbens and kalyanaraman (2012). CV is an optimal bandwidth selection procedure proposed by Ludwig and Miller (2007). The above table present estimated regression discontinuities of each interested outcome using data within specific bandwidth before and after the 3rd birthday and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). All coefficients on $Age3$ and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.

Table B.4: Donut RD for Outpatient Expenditure and Visits

	Log(outpatient expenditure)							
Size of Donut around 3rd birthday	0	3	6	9	12	15	18	21
*Age*3 (X100)	-6.90***	-6.67***	-6.84***	-6.56***	-6.20***	-6.30***	-6.61***	-6.42***
	[0.54]	[0.48]	[0.52]	[0.54]	[0.55]	[0.61]	[0.65]	[0.76]
	Log(outpatient visits)							
Size of Donut around 3rd birthday	0	3	6	9	12	15	18	21
*Age*3 (X100)	-4.73***	-4.43***	-4.42***	-4.46***	-4.37***	-4.54***	-4.70***	-4.88***
	[0.38]	[0.27]	[0.27]	[0.29]	[0.29]	[0.36]	[0.42]	[0.45]

Note: We collapse the individual-level data into age cells. Age is measured in days. The first two columns present our main results. Each observation (age cell) represents outpatient expenditures and visits from 410,517 children who were born in 2003 and 2004 (when they are age 2 and 3). Therefore, we use 2005–2008 NHI data to obtain the above estimated results. The dependent variables for the RD estimation are the log of total outpatient expenditure and the log of the total number of outpatient visits at each day of age. Each column presents estimated regression discontinuities of each interested outcome using data within 90 days before and after the 3rd birthday and report the difference in local linear regression estimates just before and after the 3rd birthday by using a triangular kernel, which gives higher weight on the data close to the 3rd birthday (equation (3.3)). we conduct a "donut" RD (Barreca et al., 2011; Shigeoka, 2014) by systematically excluding outpatient expenditure and visits within 3–21 days before and after the 3rd birthday All coefficients on *Age*3 and their standard errors have been multiplied by 100. Robust standard errors are in parentheses. *** significant at the 1 percent level, ** significant at the 5 percent level, and * significant at the 10 percent level.