Essays on Health Policy

Megan Marie Miller
Overland Park, Kansas

M.A. Economics, University of Virginia, 2016
M.S. Economics, Tufts University, 2014
B.S. Economics, University of Kansas, 2012

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of Doctor of Philosophy

Department of Economics

University of Virginia
May, 2020

In this dissertation, I study the effects of health policy. In the first chapter, I study whether households reduce their earnings to qualify for Medicaid benefits. I estimate that households are willing to reduce their earnings up to $860 to qualify for benefits and I use this to estimate the cash-equivalent value of Medicaid. I find that households would be willing to pay $614 for Medicaid benefits. The second chapter studies the relationship between Stand Your Ground laws and crime rates. Stand Your Ground laws allow individuals to use lethal force to defend themselves from danger, and I find that these laws increase crime rates. Finally, in the third chapter, I document the effects of the Medicaid expansions in 2014 on divorce rates. I find that the Medicaid expansions led to a 3% increase in the divorce rate.

Acknowledgements

I would first like to acknowledge the contributions of my committee to this dissertation. Their multitudinous comments were invaluable in refining both my approach to research and the writing of this dissertation. I would especially like to thank Leora Friedberg and John Pepper for their support and guidance at every turn.

I would also like to acknowledge the financial support of the Dean's Dissertation Completion Fellowship and the PhD+ Internship Program. Without this support, I would not have had the resources to finish this dissertation.

Finally, I would like to thank my husband, Thom, both for acting as a sounding board for my work, and for his unending support.

Contents

Chapter 1

Introduction

In this dissertation, I study the effects of health policy. In two chapters, I discuss the effect of the Medicaid program on earnings and divorces. In Chapter 2, I discuss how gun policy affects crime rates.

In the first chapter, I discuss earnings responses to Medicaid and estimate the value of benefits. Medicaid is one of the largest and costliest welfare transfer programs in the United States. Estimating the value of Medicaid provides insight on its welfare effects, but most prior work studying Medicaid has assumed this value is equal to its cost. This assumption is unlikely to hold, as factors like administrative costs and liquidity constraints may lead beneficiaries to value Medicaid at much less than its cost. To estimate this value, I study earnings responses to Medicaid using a bunching methodology that I develop for the case of a notch in the budget constraint of undetermined size. I estimate that families with earnings up to \$802 above the earnings threshold at 100% of FPL in 2004 are willing to reduce their earnings to qualify for benefits, which is a statistically significant result. I then use my estimated earnings responses and exploit variation by year and state in the location of the Medicaid earnings threshold to identify parameters of interest: the value of Medicaid, the elasticity of earnings, and adjustment costs. I estimate a value of Medicaid of \$614, which is about 24% of its cost to the government. I also estimate a statistically significant elasticity of earnings of 0.012 and that households must pay a utility cost equivalent to \$258 to adjust their earnings, and an additional \$299 to adjust earnings while maintaining benefits. My estimate of the value of Medicaid suggests that previous work has overstated the value beneficiaries place on coverage, as it is significantly less than its cost of provision.

The second chapter presents work with John Pepper on the effects of Stand Your Ground (SYG) laws. SYG laws remove the "duty to retreat" before using lethal force to defend oneself or one's property. Previous work has found conflicting estimates of the effects of these laws on crime. In this chapter, we lay out the assumptions used in previous work and how these assumptions affect the estimated relationship between SYG laws and crime rates. Additionally, we suggest a partial identification estimation strategy that weakens the necessary assumptions and provide bounds on the relationship between SYG laws and crime rates. We find small positive effects of SYG laws on homicide and violent crime rates, and uncertain effects on robbery and assault rates. In this chapter, I performed the conditional data analysis and the replications, and constructed the dataset. I also worked on the literature review and edited the chapter.

The third chapter presents evidence on the relationship between the Medicaid expansions of 2014 and divorce. In 2014, 26 states enacted the largest expansion of Medicaid in decades. One potential effect of these expansions is on divorce. The expansion of health insurance options outside of employer sponsored insurance is likely to increase divorce rates, because people that were receiving coverage through their spouse's plan now have access to coverage independent of marriage. I find no statistically significant evidence that the Medicaid expansions affect divorce rates when grouping together states that passed expansions in 2014. However, the effects of the Medicaid expansion on divorce are unlikely to be evenly distributed across the 26 states that enacted expansions, as prior to the expansion, states differed in the generosity of the program. I find statistically significant evidence that states that passed larger expansions in 2014 experienced larger increases in divorce rates. My estimates imply that the expansion increased divorce rates by around 1% on average and that the expansion increased the share of individuals divorced in the last year by around 3%.

Chapter 2

Earnings responses to in-kind transfers: The case of Medicaid

2.1 Introduction

Medicaid is one of the largest and costliest welfare programs in the United States. The goal of Medicaid is to provide health insurance to low-income Americans, and, to do this, the program enforces a strict earnings threshold. Families earning below the earnings threshold are eligible for Medicaid benefits that include no premiums and low out-of-pocket costs, while families earning above the threshold do not qualify for any benefits. This generates a “notch” in the budget constraint at the threshold, where Medicaid benefits are suddenly lost, which creates an incentive to reduce earnings to gain coverage. I estimate earnings responses to the threshold and use these responses to identify the value of Medicaid.

Most prior work has assumed that Medicaid benefits are valued at their cost. For example, Yelowitz (1998) estimates the relationship between the value of Medicaid and enrollment in Social Security Disability benefits and uses Medicaid expenditures to proxy for its value. But this assumption may not be appropriate. Administrative costs may drive up the cost of providing coverage without returning value to beneficiaries. Another possibility is that beneficiaries value Medicaid at less than its cost because they have limited resources, and would choose to take a cash transfer that is worth much less than health insurance coverage. Alternatively, Medicaid may provide access to preventive care for which beneficiaries would not be willing to pay using private insurance, but which they utilize when the care is subsidized.

Estimating the value of Medicaid provides insight on the welfare effects of a large and expensive transfer program. Previous literature on Medicaid has estimated how Medicaid benefits affect outcomes like emergency room admissions, primary care visits, bankruptcies and credit card debt.[1] These reduced form effects provide evidence that Medicaid benefits are valuable to households, but their total cash-equivalent value has been little studied. This parameter is necessary to answer questions about the efficiency of the program and to compare the relative value per dollar spent on Medicaid benefits with other welfare programs.

To estimate the value of Medicaid, I first quantify how many families locate near the notch created by the Medicaid earnings threshold. I also study how families respond to changes in the location of these notches. I focus on the six states with a Medicaid threshold at 100% of the federal poverty line (FPL) in 2004.[2] Of these states, three moved the Medicaid threshold in 2005, providing an opportunity to observe bunching before and after the earnings threshold changes location.

I estimate that families with earnings up to \$802 above the earnings threshold at 100% of FPL in 2004 are willing to reduce their earnings to qualify for benefits. When the threshold is lowered, this earnings response falls to \$427 and when the threshold is raised, the earnings response is \$744. This means that households are more willing to adjust their earnings in response to a decrease in the threshold than in response to an increase in the threshold. I estimate these responses using Survey of Income and Program Participation (SIPP) and Current Population Survey (CPS) data from 2004-2005. I also estimate earnings responses after accounting for rounding in self-reported earnings, and find similar results.

Next, I use my estimated earnings responses and exploit variation by year and state in the location of the Medicaid earnings threshold to identify my parameters of interest: the value of Medicaid, the elasticity of earnings, and adjustment costs. I construct a model of labor supply and match my estimated observed earnings responses with predicted earnings responses from the model. I include adjustment costs in my model, so that households must pay a fixed utility cost to adjust their earnings. These adjustments costs could take

[1](Taubman et al 2014), (Curry and Gruber 2000),(Gross and Notowidigdo 2011), (Hu et al 2016)

[2]I focus on 100% of FPL because it was the Medicaid earnings threshold that was shared by the largest number of states in 2004.These states are: California, Connecticut, Delaware, Ohio, Oregon, and Tennessee. Connecticut, Ohio, and Tennessee changed their thresholds in 2005. Details on the Medicaid program and earnings thresholds are provided in the next section

the form of traditional frictions that prevent workers from changing their work hours or place of employment to target a particular level of earnings, or they could take the form of frictions related to maintaining Medicaid coverage.

The direction of the change in the earnings threshold identifies my parameters of interest. When the Medicaid earnings threshold increases, no workers who located at the initial threshold choose to leave the program, but they may increase their earnings and maintain eligibility. This choice depends on the size of adjustment costs. When the Medicaid earnings threshold decreases, some workers will leave Medicaid instead of reducing their earnings further to maintain eligibility. Their choice to leave depends both on the size of adjustment costs and the value of Medicaid. This variation allows me to identify the value of Medicaid to beneficiaries, along with adjustment costs.

I estimate a value of Medicaid of \$614, which is about 24% of its cost to the government. My estimate is statistically significant, with a standard error of \$124. I also estimate an elasticity of earnings of 0.012, which implies that low-income workers do not respond to changes in marginal tax rates. This estimate is also statistically significant. Both estimates are similar to those in the literature. Households also face large costs to adjusting their earnings. My estimate implies that, in order to adjust their earnings, households must pay a utility cost equivalent to \$258. There is an additional cost of \$299 that affects households adjusting earnings while maintaining benefits, which may reflect the cost of verifying earnings through the program. My estimates of adjustment costs are noisier, and are insignificant, with large standard errors. My results are similar across the SIPP and CPS, and are in line with estimates of the value of Medicaid in the literature.

My estimates of the earnings response are generally small, and in some cases statistically insignificant. There is a sizable literature that studies the labor supply effects of the Medicaid program. This literature generally finds that Medicaid has a limited effect on labor supply. Yelowitz (1995) uses Medicaid expansions that detach Medicaid eligibility from cash assistance eligibility to show that the labor supply effects of public insurance coverage are large for married women, and insignificant for unmarried women. Meyer and Rosenbaum (1999) also study Medicaid expansions of the same time period and find similarly limited labor supply responses of single women to Medicaid. Finally, Baicker et al (2014) find small, negative, and statistically insignificant labor supply responses to the Oregon Health Insurance Experiment.

I contribute to the literature using non-linear budget set methods to estimate labor responses to kinks in the budget constraint, where the marginal tax rate changes abruptly along the budget constraint, and notches caused by tax policies.[3] I adapt these methods to the case of an in-kind transfer that creates a notch in the budget constraint of unknown size. Notches change the average tax rate along the budget constraint by introducing a level drop in after-tax earnings, or consumption, at the earnings threshold. Notches created by in-kind transfer programs are of unknown size, because the drop in consumption associated with the earnings threshold depends on the cash-equivalent value of the transfer.

I also contribute estimates of the value of Medicaid coverage, which is a parameter that has received little attention in the literature. Two papers that estimate this value are Finkelstein, Hendren, and Luttmer (forthcoming) and Gallen (2015). These papers find that the value of Medicaid is around 25%-50% of its cost of provision. This is similar to my estimate that Medicaid is valued at around 24% of its cost. My estimation strategy does not require detailed Medicaid administrative data, as Finkelstein, Hendren, and Luttmer (forthcoming) use, and does not rest on a structural model of labor supply, which Gallen (2015) estimates.

In Sections (2.2) and (2.3), I describe the Medicaid program and my data, respectively. I present my model of labor supply in Section (2.4) and in Section (2.5), I outline my empirical strategy. Section (2.6) contains my estimates of labor supply responses to Medicaid and key parameters. Section (2.7) concludes.

2.2 Medicaid Description

During the period I consider, in the mid-2000s, each state set its own Medicaid eligibility criteria and coverage details, subject to federal requirements. Each state chose an earnings threshold, above which households did not qualify for coverage, and an asset test, which disqualified households with significant assets from receiving benefits. I discuss each of these details below.

One component of the eligibility criteria was an earnings threshold. The relevant measure of income during this period for Medicaid eligibility was earned income, and the threshold was expressed as a percentage of the federal poverty line. For a family of three in 2004, the federal poverty line was set at $15,670. The federal poverty line differs by family

[3]Kleven and Waseem (2013); Gelber, Jones and Sacks (2014); Saez (2010)

size, which means that the location of the Medicaid earnings threshold depends both on family size and state of residence.

In the mid-2000's, Medicaid was intended to cover low-income children, pregnant women, and the disabled. However, states could offer coverage that extended beyond federal guidelines and set less restrictive eligibility thresholds for parents in low-income families. This meant that during this period, there was significant variation across states in the location of the Medicaid earnings threshold for parents. In 2004, 26 states offered coverage to parents only in very low-income families, earning between 20% and 75% of FPL. The remaining states offered coverage to parents earning near or above the poverty line.

I study the six states with a Medicaid earnings threshold for parents coverage at 100% of FPL. This is primarily because survey data provides self-reported income, meaning that income is noisy in my data. That makes it difficult to observe excess mass in a small sample. To ameliorate this problem, I pool across the states that share an eligibility threshold. In 2004, 100% of FPL was the Medicaid earnings threshold that was shared by the largest number of states. I do not study bunching in states with earnings thresholds at locations other than 100% of FPL because my parameter estimates are drawn from households located in the bunching interval. Households with earnings near 150% of FPL are unlikely to have the same elasticity, adjustment costs, or value of Medicaid as individuals with earnings at 100% of FPL.

The six states that I focus on are: California, Connecticut, Delaware, Ohio, Oregon, and Tennessee. None of these states moved their thresholds between 2003 and 2004, while three of them moved their earnings thresholds in 2005. Ohio decreased its earnings threshold to 90% of FPL, while Tennessee decreased its earnings threshold to 81% of FPL. Connecticut increased its earnings threshold to 150% of FPL. These movements in the earnings thresholds will serve to identify the elasticity of earnings with respect to the marginal tax rate, adjustment costs, and the value of Medicaid.

In 2004, a little over 55 million Americans were covered by Medicaid. Of these 55 million, around 48% were children, 22% were adults, 8% were aged, and 14% were disabled. While aged and disabled individuals make up a minority of beneficiaries, they represent the largest share of expenditures on healthcare by the Medicaid program. In 2004, the program spent $1,671 per child served, $2,509 per adult and $14,070 per disabled enrollee.[4]

I focus on adult, non-disabled enrollees, who are the most likely to adjust earnings to gain coverage. Disabled beneficiaries often qualify for Medicaid through Social Security disability coverage, which limits the number of hours a beneficiary is allowed to work. Aged beneficiaries are Medicaid recipients who also qualify for Medicare, meaning they are over 65 and likely to be retired. As a result, disabled and aged Medicaid recipients are unlikely to adjust earnings in response to changes in the threshold location.

To the extent that we observe discontinuities in the earnings distribution near 100% of FPL, we need to be certain that these are related to the Medicaid notch, and not to another welfare program or tax policy. Figure (2.1) depicts the location of transfer program earnings tests near 100% of FPL. From the figure, we can see that the kink in the budget constraint created by the Food Stamp program is not located near the Medicaid notch, while the notches created by Children's Health Insurance Programs (CHIP/SCHIP) are, in some cases, located near the Medicaid notch. A kink point in the earnings distribution created by the Earned Income Tax Credit was also located at $15,000 in 2004 for married taxpayers. In the results section, I perform robustness checks to ensure that the results from my estimation strategy are caused by the Medicaid notch, and not another transfer program with an earnings test near 100% of FPL or the EITC kink point.

There are a couple institutional details regarding Medicaid that I cannot address with my empirical strategy, but which may affect the analysis. First, some states imposed an asset test in determining eligibility for benefits. Each state determined which assets counted towards the asset test and the level at which applicants were ineligible for benefits. In the mid-2000's, many states were removing these tests, but they were still common for parent coverage and nearly nonexistent for children's coverage. In 2004, 21 states had no asset test for parent coverage. For the six states I consider in my analysis, California, Tennessee, and Oregon enforced an asset test in 2004 (Kaiser Family Foundation 2004).

Without accounting for these asset tests, I might underestimate bunching. I would conclude that households earning near the earnings threshold are unresponsive to Medicaid because they do not value the benefit, when in fact they may not be eligible due to assets. A 2001 Kaiser Family Foundation report showed that, when Louisiana removed their asset test for parent's coverage, they experienced a 3% increase in enrollment, indicating that the

[4]Program details come from the Center for Medicare and Medicaid Services *2012 Statistical Supplement*

asset test was not a significant barrier to enrollment.

Finally, each state also determined the frequency at which earnings were verified and the number of months of earnings that were considered in determining eligibility. Beneficiaries generally verified their earnings every 6-12 months, and monthly earnings are considered. In 2004, 37 states had a 12 month renewal period for parent coverage, with some states verifying income more frequently. These differences would add noise to my estimates of bunching. This would mean that Medicaid recipients may report higher monthly earnings in some months, when their earnings will not be verified for Medicaid eligibility. This would introduce a form of measurement error in my estimates of bunching, as some households may have reduced their earnings to qualify and then increased their earnings, meaning that observed bunching would not be sharp at the earnings threshold when annual measures of earnings are used.

2.3 Data Description

Because estimates of bunching require a large number of observations, I consider several sources of data, which I describe in this section. The two datasets I use in my analysis are the Survey of Income and Program Participation (SIPP) and the Current Population Survey (CPS).

The 2004 SIPP is a quarterly panel survey of 51,379 households that asks respondents questions about sources of income and participation in transfer programs in the United States. The SIPP is typically a four-year panel, but the 2004 SIPP was cut short at 2 1/2 years and covers 2004 through mid-2006. I use data from 2004 and 2005, the only two complete years in this panel.

The location of the Medicaid threshold varies by family size and state, so I consider 3-person households in states with a Medicaid earnings threshold at 100% of FPL. I choose 3-person households because they are the most common family size with earnings near the Medicaid earnings thresholds. To show how this affects my sample size, I have included household and individual observation counts in Table (2.1). The number of individuals is higher than the number of households because both surveys collect respondent-level data. First, around 17% of families have 3-members, which brings the household count in the SIPP to 8,337. Next, around 14% of households in the SIPP live in the six states I consider, which means that my estimates of bunching are estimated off of 1,469 households.

Restricting my analysis to a handful of states and 3 member families reduces my sample size significantly, and so I also use CPS data, which is a larger sample. In 2004, the CPS surveyed 76,439 families, and 2,318 of these families had 3 members in the six states I consider. This data represents a repeated cross-section rather than a panel. The benefit of using panel data, like the SIPP, is that any observed changes in the earnings distribution must be caused by people who adjust their earnings, while in a repeated cross-section, observed changes in the earnings distribution may be caused by changes in the sample. To investigate this possibility, I estimate bunching using the SIPP and the CPS separately to see whether I observe different patterns between the two data sources.

The CPS March Supplement asks about annual sources of income, and I construct the measure of annual earnings that is relevant for Medicaid eligibility, which is household earnings that does not include government transfers. The SIPP asks about monthly sources of income, which I aggregate to construct a measure of annual income. By using annual data, I can combine CPS and SIPP earnings data.

In the second panel of Table (2.1), I present summary statistics for 3 member households to understand how these households differ from others. For comparison, summary statistics for all households in both surveys are presented in the first panel of Table (2.1). In both the SIPP and CPS, 3 member households are higher-earning than average, more likely to be married, and more likely to have young children. The adults in these households are also younger on average and slightly more likely to work. I also present summary statistics for 3 member households in the six states on which I focus in the third panel of Table (2.1), and these households do not appear to be different from all 3 member households in either survey.

I have also included summary statistics for both the 2004 SIPP and 2004 CPS in Table (2.1), to look for important differences in the two surveys. Because I am using only 3 member households in the six states with a Medicaid threshold at 100% of FPL, the relevant summary statistics for this comparison are in the last panel of Table (2.1). The most notable difference between the surveys is the lower average household income among SIPP respondents. The SIPP intentionally over-samples low income households, because one purpose of the survey is to understand participation in welfare programs. SIPP respondents are also slightly more likely to be married and slightly more likely to work. Three member households in the SIPP are also more likely to have children under 6.[5]

In Figure (2.2), I've included a graph of the share of individuals in the SIPP and CPS receiving Medicaid benefits across the earnings distribution in the 6 states with a Medicaid earnings threshold at 100% of FPL. There is a pretty clear drop near $16,000, but there are still some people that report receiving benefits while earning above the threshold. This could either be because families increase their earnings after receiving benefits, as explained above, or because they do not know if their benefits are from CHIP/SCHIP or a state program.

2.4 Model

The model of notches below builds on Saez (2010) and Kleven and Waseem (2013). Kleven and Waseem (2013) study notches in the budget constraint created by the Pakistani tax system, and identify the elasticity of taxable income and earnings frictions. I adapt this model to the case of Medicaid, which creates a notch in the budget constraint that is of unknown size.

This model shows how bunching, or excess mass in the observed earnings distribution, can be used to identify the elasticity of earnings with respect to the marginal tax rate, the cost of adjusting earnings, and the size of the Medicaid notch. I lay out the model in two sections. In the first section, I present a one-period model to show how to express bunching as a function of the elasticity of earnings with respect to the marginal tax rate. The one period model that I present below follows Kleven and Waseem (2013) closely. In the second section, I present a two-period model to show how movements in the location of the notch can be used to identify key parameters.

2.4.1 One period model

In this model, there are two goods: before-tax earnings, z, and after-tax earnings, c. In order to earn income z, individuals must work and incur some utility cost in the form of

[5] Among these six states, 48% live in Ohio, Tennessee, and Connecticut, the three states that moved their Medicaid earnings threshold in 2005. The median family size among all families in the survey is 2, and 17% of families in this survey have three members. The sample loss rate for the 2004 SIPP was 31.2% after 8 waves, or two years.

effort. The utility function is defined as $u(c, z, n)$,where n is a random variable reflecting heterogeneity in the utility cost of effort. The individual faces a marginal tax rate, t, and there is a notch in the budget constraint of size T, located at earnings level z^*. The budget constraint for this individual is depicted in Figure (2.3). In this case, T is the value of the in-kind transfer that is available as long as earnings are below z^*.

I also allow for the existence of adjustment costs. If adjustment costs are not included in the model, the elasticity of earnings and value of Medicaid would be underestimated. Adjustment costs prevent households from responding to the Medicaid notch, and this lack of response would be attributed either to a low elasticity of earnings or a low value of Medicaid if adjustment costs are not included in the model.

Figure (2.4) depicts the indifference curves of two households with utility functions, U and V. Household V would choose earnings z_V if there were no notch at z^*. This household receives utility V_A at the notch, and utility V_B at z_V. Without adjustment costs, this household would prefer to move to the notch, because $V_A > V_B$. However, they stand to gain little by moving to the notch, so that $V_A - V_B < a$, where a is the utility cost of adjusting earnings. In other words, the benefit of moving to the notch is less than the cost incurred. This household will choose not to move to the notch if they face adjustment costs of a.

On the other hand, household U in Figure (2.4) is exactly indifferent between remaining at their optimal earnings in absence of the notch, $z^* + \Delta z_a^*$, and paying an adjustment cost and moving to the notch at z^*. U_A is their utility at the notch, while U_B is their utility at $z^* + \Delta z_a^*$. The indifference condition implies that $U_A - U_B = a$.

Households who would choose earnings less than $z^* + \Delta z_a^*$ in the absence of the notch experience a larger utility gain from moving to the notch, so they adjust their earnings, while those who would choose earnings greater than $z^* + \Delta z_a^*$ in the absence of the notch will prefer not to move. The earnings response to the notch is defined as Δz_a^*, or the response to the Medicaid notch of the indifferent household.

Bunching is defined as excess mass in the earnings distribution caused by the notch. This excess mass comprises all households that move to the notch. Thus, bunching is the mass of households who would earn between z^* and $z^* + \Delta z_a^*$ in absence of the notch, and is a function of Δz_a^*, the earnings response.

This graphical analysis describes the relationship between bunching and the earnings response Δz_a^*. In order to solve for an expression for bunching in terms of the structural elasticity, e, and adjustment cost, a, we need to make an assumption about the utility function. Following the literature, suppose the utility function is given by $u(c, z, n) = c - \frac{n}{1+\frac{1}{e}} \left(\frac{z}{n}\right)^{1+\frac{1}{e}}$, where $c = z(1-t)$. The structural elasticity is e, and n is an ability parameter with a smooth, continuous density, $f(n)$, and cumulative distribution $F(n)$.

Optimal earnings in the absence of the notch, which I call counterfactual earnings going forward, are given by the solution to the utility maximization problem:

$$\underset{c,z}{\text{maximize}} \quad c - \frac{n}{1+\frac{1}{e}} \left(\frac{z}{n}\right)^{1+\frac{1}{e}} \quad \text{subject to } c = (1-t)z$$

Solving this problem yields an expression for counterfactual earnings, $z = n(1-t)^e$. Counterfactual earnings are a function of the random ability parameter, n, the marginal tax rate, t, and the elasticity of earnings, e. The counterfactual earnings distribution, denoted $H_0(z)$, can be written as $H_0(z) = H_0(n(1-t)^e) = P(n(1-t)^e < z) = F\left(\frac{z}{(1-t)^e}\right)$. Because the distribution of n, the ability parameter, is smooth and continuous, $h_0()$, the counterfactual earnings distribution, is as well.

The indifferent household, with counterfactual earnings $z^* + \Delta z_a^*$, is indifferent between earning their counterfactual earnings level, and paying the adjustment cost, a, and moving to the notch. From Figure (2.4), the indifference condition implies that $U_A - U_B = a$, or:

$$u((1-t)z^*, z^*, n^* + \Delta n_a^*) - u((1-t)(z^* + \Delta z_a^*) - T, z^* + \Delta z_a^*, n^* + \Delta n_a^*) = a \quad (2.1)$$

The ability level of the indifferent household is $n^* + \Delta n_a^*$. Using our solution for counterfactual earnings, we can write $n^* + \Delta n_a^* = (z^* + \Delta z_a^*)(1-t)^{-e}$. Equation (2.1) is then an implicit function that relates the earnings response, Δz_a^*, to adjustment costs and the elasticity.

Bunching is the mass of households with counterfactuals earnings between the notch at z^*, and the counterfactual earnings of the indifferent household, $z^* + \Delta z_a^*$. Now that we can express the earnings response in terms of adjustment costs and the structural elasticity, we can express the level of bunching in terms of these two parameters.

$$B = \int_{z^*}^{z^* + \Delta z_a^*} h_0(z)dz \approx h_0(z^*)\Delta z_a^* \quad (2.2)$$

The approximation assumes that the counterfactual earnings distribution is constant over the bunching interval. Equation (2.2) expresses bunching in terms of the earnings response. From equation (2.1), we know that the earnings response is in turn an implicit function of the elasticity of earnings and adjustment costs. Therefore, the level of bunching is determined by the elasticity of earnings with respect to the tax rate, e, and adjustment costs, a.

2.4.2 Two period model

In the previous section, we expressed bunching as a function of two unknown parameters, the elasticity of earnings, e, and adjustment costs, a. One additional layer of complexity added by considering the Medicaid notch is that the size of the notch, T, which represents the cash equivalent value of the transfer, is unknown. A contribution of this paper is to estimate the cash-equivalent value of Medicaid for families near the earnings threshold.

To consider a model in which the value of Medicaid, T, is unknown, I introduce a two period model where the notch exists in the first period, and moves in the second period. I evaluate bunching at the initial notch after the notch moves, both in the case where the notch increases and where it decreases. This two-period model builds on Gelber, Jones and Sacks (2013), who use bunching at kinks in the earnings distribution caused by the Social Security earnings test to estimate the elasticity of taxable income. They also use residual bunching after the location of the kinks has shifted to identify adjustment costs.

In the first period, I assume that households face no cost to adjust their earnings and move to the initial notch at z^*. The relevant indifference condition is then $U_A - U_B = 0$, or:

$$u((1-t)z^*, z^*, n^* + \Delta n_a^*) - u((1-t)(z^* + \Delta z_a^*) - T, z^* + \Delta z_a^*, n^* + \Delta n_a^*) = 0 \quad (2.3)$$

Now consider the case where the earnings threshold increases, which is shown in Figure (2.5). In Figure (2.5), the threshold moves from z^* to $\tilde{z}_I$. When the earnings threshold increases to $\tilde{z}_I$, households who previously qualified for Medicaid can now choose to increase their earnings while maintaining coverage. In order to adjust their earnings while maintaining eligibility for the transfer, households must pay an adjustment cost of a_1. Figure (2.5) depicts the household who is just indifferent between moving from the initial notch, at z^*, to their counterfactual earnings, at $\bar{z}_I$. This household faces the indifference condition $U_G - U_H = a_1$, which is given by:

$$u((1-t)\bar{z}_I, \bar{z}_I, \bar{n}_I) - u((1-t)z^*, z^*, \bar{n}_I) = a_1 \tag{2.4}$$

Combined with the solution for counterfactual earnings, $\bar{n}_I = \bar{z}_I(1-t)^{-e}$, the indifference condition expresses the counterfactual earnings of the indifferent household, $\bar{z}_I$, as an implicit function of e and a_1.

To characterize bunching at z^* when the threshold moves, consider households with counterfactual earnings greater than $\bar{z}_I$. These households gain more than the indifferent household by moving away from the initial notch, and will not continue to bunch at z^*. Meanwhile, anyone with counterfactual earnings less than $\bar{z}_I$ gains less than the indifferent household by moving, and will prefer to remain at the initial notch. So, bunching at the initial notch after the move is given by:

$$B_{2I} = \int_{z^*}^{\bar{z}_I} h_0(z)dz \approx h_0(z^*)(\bar{z}_I - z^*) \tag{2.5}$$

Bunching at the initial notch after the notch increases, B_{2I}, is a function of $\bar{z}_I$, and so is also a function of e and a_1.

Next consider what happens when the threshold moves down. At the initial notch, now households lose the value of the notch, T. Households that were initially located at z^* will choose either to remain at z^*, move to the new notch, or move to their counterfactual earnings. I allow for the possibility that adjustment costs are asymmetric, so that it is more costly to adjust earnings while maintaining Medicaid coverage than to adjust earnings outside of Medicaid coverage. This is both to better match observed patterns of bunching and to reflect the reality that verifying earnings for Medicaid eligibility can be burdensome. To adjust earnings while maintaining coverage, households must pay a_1, while the cost to adjust earnings while not receiving benefits is a_2. To characterize the decision to remain at the initial notch, consider Figure (2.6). Indifference curves for two households, U and V, are depicted.

Suppose the new notch is located at $\tilde{z}_D < z^*$. In Figure (2.6a), household U, with counterfactual earnings $\bar{z}_D^L$, is just indifferent between remaining at the initial notch, at z^*, and moving to the new notch, at $\tilde{z}_D$. U_J is U's utility at the new notch, while U_K is their utility through the initial notch. Household U faces the indifference condition $U_G - U_H = a_1$. Households with counterfactual earnings less than $\bar{z}_D^L$ gain more by moving to the new notch, and so will not continue to bunch at the initial notch.

Household V in Figure (2.6b) has counterfactual earnings $\bar{z}_D^U$. Household V is just indifferent between moving from the initial notch, at z^*, to their counterfactual earnings, $\bar{z}_D^U$. V_J is their indifference curve through z^*, after the notch moves, while V_K is their indifference curve through counterfactual earnings $\bar{z}_D^U$. The indifference condition implies that $V_J - V_K = a_2$. Everyone with counterfactual earnings above $\bar{z}_D^U$ will move from the initial notch to their counterfactual earnings.

Plugging in counterfactual earnings, the indifference conditions that define $\bar{z}_D^L$ and $\bar{z}_D^U$ are:

$$u(\bar{z}_D^U(1-t) - T, \bar{z}_D^U, \bar{n}_D^U) - u((1-t)z^* - T, z^*, \bar{n}_D^U) = a_2 \tag{2.6}$$

$$u(\tilde{z}_D(1-t), \tilde{z}_D, \bar{n}_D^L) - u((1-t)z^* - T, z^*, \bar{n}_D^L) = a_1 \tag{2.7}$$

These two conditions express $\bar{z}_D^U$ and $\bar{z}_D^L$ as implicit functions of e, a_1, a_2, and T. Bunching at the initial notch, when the earnings threshold falls is then the mass of all households with counterfactual earnings between $\bar{z}_D^L$ and $\bar{z}_D^U$. This is given by:

$$B_{2D} = \int_{\bar{z}_D^L}^{\bar{z}_D^U} h_0(z)dz \approx h_0(z^*)(\bar{z}_D^U - \bar{z}_D^L) \tag{2.8}$$

Note that for some parameter values, there may be no household who was willing to bunch both at z^* and at $\tilde{z}_D$. In this case, $\bar{z}_D^L$ is equal to z^*, and bunching at the initial notch when the threshold falls is determined only by the household indifferent between remaining at the initial notch and increasing their earnings to their counterfactual earnings.

Finally, consider what happens at the new earnings threshold when the threshold falls. A decline in bunching at z^* when the threshold falls could be due to households reducing their earnings to the new notch, because the transfer is valuable, or to households increasing their earnings to their counterfactual earnings because the transfer is not valuable. In this way, bunching at the new threshold is determined by the value of the transfer, because the decision of households to either decrease their earnings to the new threshold or increase their earnings to their counterfactual earnings level depends on the value of the transfer.[6]

When the threshold falls, the indifference condition that determines the counterfactual earnings of the indifferent household is given by:

$$u(\tilde{z}_D(1-t), \tilde{z}_D, \bar{n}_{2D}) - u((1-t)\bar{z}_{2D} - T, \bar{z}_{2D}, \bar{n}_{2D}) = a_1 \tag{2.9}$$

This indifference condition holds as long as $\bar{z}_D^L$ is equal to z^*. In other words, as long as the indifferent household when the threshold falls was not willing to bunch when the earnings threshold was at z^*.

Going forward, it will be useful to define notation for bunching weighted by the counterfactual distribution at the notch.

$$b = \frac{B}{h_0(z^*)} \tag{2.10}$$

This expresses bunching as how many more households are located near the notch than we would expect if there were no discontinuity in the budget constraint. This is the expression that I call "predicted" bunching from the model in later sections.

2.5 Empirical Strategy

The model section shows how to express predicted bunching in terms of my parameters of interest. In the empirical strategy section, I first discuss how I estimate observed bunching near the Medicaid notch using a discontinuity design. Then, I match the estimated bunching to predicted bunching from the model by choosing e, a_1, a_2, and T.

2.5.1 Bunching

I follow Gelber, Jones, and Sacks (2014) and Kleven and Waseem (2013) to estimate bunching near the Medicaid notch at 100% of FPL. The method involves fitting a polynomial through observed earnings, excluding the contribution of observations near the threshold.

[6] I choose not to study bunching at the new threshold when the threshold increases, because, in my setting, this would mean that my parameter estimates are drawn from the population of workers with earnings between around \$13,000 and \$21,000. By using only bunching when the threshold falls, my estimates are drawn only from households with earnings near \$13,000 and \$16,000.

This is the estimated counterfactual earnings distribution. Bunching is then measured as the excess mass located near the threshold, above the mass predicted by the counterfactual earnings distribution. I estimate the following regression:

$$p_i = \sum_{d=0}^{D} \beta_d (z_i - z^*)^d + \sum_{j=-k}^{j=k} \gamma_j * 1(z_i - z^* = j\delta) + u_i \tag{2.11}$$

Here, p_i is the proportion of observations in the interval $(z_i \pm \delta/2)$ and z^* is the notch location. Estimates of the counterfactual density at each bin z_i are given by $\sum_{d=0}^{D} \beta_d (z_i - z^*)^d$. D is the degree of the polynomial that is estimated for the counterfactual earnings distribution, and δ is the size of the bins into which earnings are grouped. The number of observations is equal to the number of bins, which is 235 for the main specification. This means that I estimate the earnings distribution up to \$190,000, at which point the density of observations in each bin is near zero. I set $D = 7$, and $\delta = 800$.

Bunching is a measure of how many more households locate near the notch than would be expected if there were no discontinuity in the budget constraint, and is estimated as:

$$b = \frac{\sum_{j=-k}^{j=k} \hat{\gamma}_j * 1(z_j - z^* = j\delta)}{\sum_{j=-k}^{j=k} \sum_{d=0}^{d=D} (j\delta)^d} \tag{2.12}$$

In words, k is the number of excluded bins on either side of the notch, and I set $k = 2$. [7] This means that bunching is the excess mass located within \$2,000 of the notch. To see this, note that the bin containing the Medicaid notch runs from \$15,270 to \$16,070. There are 2 bins on either side of the central bin, which means that I include indicators for bins running between \$13,670 and \$17,670. In the appendix, I present results from estimating parameters using alternative parameter choices.

[7] Standard errors are calculated using the method outlined in Chetty, Friedman, Olsen, and Pistaferri (2011). First, equation (9) is run, and coefficients estimates are obtained. Then, 50 random samples of size 235 are taken from the residuals and used to generate new observations of the proportion of observations in an earnings bin, the dependent variable. The regression is run again on these simulated observations, and the standard error of the coefficient estimate is the standard error from the 50 estimated coefficients using the simulated observations.

2.5.2 Round number bunching

I also consider how estimates of bunching are affected by round number bunching. Round number bunching creates discontinuities in the counterfactual earnings distribution, because some people report their earnings as round number multiples of $5,000, which causes peaks in the observed earnings density at these points. This would mean that a counterfactual earnings distribution estimated as a smooth polynomial would underestimate the fraction of earnings reported as a multiple of $5,000. My estimates are particularly sensitive to this issue, because I exclude earnings near $15,000 to calculate excess mass, so people that rounded their earnings to $15,000 are included in my estimate of bunching. To control for this, I include a dummy variable for bins that include multiples of $5,000. I estimate the following regression:

$$p_i = \sum_{d=0}^{D} \beta_d (z_i - z^*)^d + \sum_{j=-k}^{j=k} \gamma_j * 1(z_i - z^* = j\delta) + \alpha * I_i + u_i \tag{2.13}$$

Here, I_i is an indicator variable that equals one if a bin contains a multiple of $5,000. Then, bunching is the fraction of people reporting earnings in the excluded bins, or $\sum \gamma_j$, after accounting for the expected share predicted by the seventh order polynomial and round number bunching at multiples of $5,000.

2.5.3 Elasticity, adjustment costs, and the value of Medicaid

I next estimate the elasticity of earnings, adjustment costs, and the value of Medicaid using bunching at the initial notch, before and after the notch moves location. This is similar to Gelber, Jones, and Sacks (2014). As discussed in the model section, bunching before and after the notch moves location are each a function of the unknown parameters. So, to estimate them, I minimize the difference between observed and predicted bunching by choosing elasticity, e, the cost of adjusting earnings while maintaining Medicaid coverage, a_1, and the cost of adjusting earnings while not receiving benefits, a_2, and the value of Medicaid, T.

In particular, I choose (e, a_1, a_2, T) to solve the following minimization problem:

$$\underset{e,a_1,a_2,T}{\text{minimize}} \quad (\hat{b}(e, a_1, a_2, T) - b)' W' (\hat{b}(e, a_1, a_2, T) - b)$$

Here, b is a 4x1 vector of observed bunching in both periods, as calculated using equation (2.12). $\hat{b}$ is predicted bunching from the model, given e, a_1, a_2 and T. W' is a 4x4 weighting matrix with the inverse of the standard error of each of the observed bunching estimates.

To calculate predicted bunching from the model, I use the indifference condition, equation (2.3), and solve numerically for the earnings response, Δz_a^*. Recall that equation (2.3) is the indifference condition for the household that is just indifferent between bunching at the notch, z^*, and earning $z^* + \Delta z_a^*$. After some simple algebra, and plugging in to the utility function, I numerically solve the following equation for Δz_a^*, given a guess for e, a_1, a_2 and T:

$$\frac{1}{1+\frac{\Delta z_a^*}{z^*}}\left(1+\frac{T}{z^*}\frac{1}{1-t}\right)-\frac{1}{1+\frac{1}{e}}\left(\frac{1}{1+\frac{\Delta z_a^*}{z^*}}\right)^{1+\frac{1}{e}}-\frac{1}{1+e}=0 \tag{2.14}$$

Then, from equations (2.2) and (2.10), predicted bunching in the first period, $\hat{b}_1$, is given by:

$$\hat{b} = \Delta z_a^*$$

To evaluate predicted bunching at z^* after the notch moves, denoted $\hat{b}_2$, I need to consider how bunching changes in states where the threshold increases, and states where the threshold decreases.

First, we will consider predicted bunching at z^* when the threshold increases. From equation (2.4), the indifference condition in Connecticut, where the Medicaid earnings threshold increased, is:

$$(1-t)(z^* - \tilde{z}_I) + a_1 = (1-t)\left(\frac{\bar{z}_I^{-\frac{1}{e}}}{1+\frac{1}{e}}\right)\left[{z^*}^{1+\frac{1}{e}} - \tilde{z}_I^{1+\frac{1}{e}}\right]$$

Where $\tilde{z}_I$ represents the new threshold and z^* is the initial threshold. With e, a_1 and T, we can solve for $\bar{z}_I$, or the counterfactual earnings of the household that is just indifferent between remaining at z^* and moving to the new notch at $\tilde{z}_I$. From equation (2.5), predicted bunching at z^* after the notch moves, among households in states where the Medicaid notch increases is given by:

$$\hat{b}_{2I} = \bar{z}_I - z^*$$

In states where the Medicaid notch moves down, there are two indifference conditions that determine both the upper and lower bound for bunching. As in the model, these thresholds are denoted by $\bar{z}_D^U$ and $\bar{z}_D^L$. The new notch location is denoted $\tilde{z}_D$. With a guess for e, a_1, a_2, and T, we can solve for these two thresholds using equations (2.6) and (2.7) from the model:

$$(1-t)(\tilde{z}_D - z^*) + T - a_1 = (1-t)\left(\frac{(\bar{z}_D^L)^{-\frac{1}{e}}}{1+\frac{1}{e}}\right)\left[\tilde{z}_D^{1+\frac{1}{e}} - (z^*)^{1+\frac{1}{e}}\right]$$

$$(1-t)(\bar{z}_D^U - z^*) - a_2 = (1-t)\left(\frac{(\bar{z}_D^U)^{-\frac{1}{e}}}{1+\frac{1}{e}}\right)\left[(\bar{z}_D^U)^{1+\frac{1}{e}} - (z^*)^{1+\frac{1}{e}}\right]$$

From equation (2.8), predicted bunching in the second period in states where the threshold moves down is:

$$\hat{b}_{2D} = \bar{z}_D^U - \bar{z}_D^L$$

As mentioned in the model section, for some values of e, a_1, T, and a_2, there will be no households that bunch at both the initial threshold and the new threshold when the threshold moves down. In this case, $\bar{z}_D^L = z^*$.

Finally, predicted bunching in the second period at the new threshold, $\tilde{z}_D$ is given by:

$$(1-t)(\tilde{z}_D - \bar{z}_{2D}) + T - a_1 = (1-t)\left(\frac{(\bar{z}_{2D})^{-\frac{1}{e}}}{1+\frac{1}{e}}\right)\left[\tilde{z}_D^{1+\frac{1}{e}} - (\bar{z}_{2D})^{1+\frac{1}{e}}\right]$$

2.5.4 Identification

It is helpful to summarize the intuition behind identification in this model. As in the Saez (2010) model, the elasticity of earnings is identified by bunching in the initial period. The asymmetric cost to adjust earnings while maintaining Medicaid eligibility is identified by bunching at the initial threshold, z^*, after the threshold increases. This is because, when the threshold increases, workers earning near z^* can increase their earnings while maintaining Medicaid coverage, and so any workers that continue to earn z^* must be facing earnings frictions while maintaining eligibility.

Earnings frictions caused by traditional adjustment costs are identified by bunching at the initial threshold when the threshold falls. This is because these workers no longer receive Medicaid coverage at the initial threshold. Whether these workers choose to move away from z^* identifies the earnings frictions affecting workers that are no longer receiving coverage. Allowing for these asymmetric costs allows the model to explain why bunching at the initial threshold is lower when the threshold falls than when it increases. Finally, bunching at the new threshold, $\tilde{z}_D$, identifies the value of Medicaid. The extent to which individuals are willing to reduce their earnings, pay all adjustment costs, and bunch at the new threshold identifies the value of the transfer.

2.6 Empirical Results

I present my empirical results in three parts. First, I show graphical evidence of bunching and adjustment costs in earnings distributions. I also present evidence that the EITC and CHIP thresholds are not confounding my bunching estimates. Second, I discuss the results of estimating the counterfactual earnings distribution by running equation (2.11) and equation (2.13) and I quantify bunching near the Medicaid earnings threshold using equation (2.12). To provide intuition about the relationship between the elasticity of earnings and the value of the transfer, I also use my estimate of bunching to calculate an approximation of the elasticity of earnings. Finally, I present estimates of my four parameters of interest: the elasticity of earnings, e, adjustment costs, a_2, the additional cost to adjust earnings while receiving the transfer, $a_1 - a_2$, and the value of the transfer, T.

2.6.1 Graphical evidence of bunching

To motivate my estimation strategy, I present kernel density estimates that are calculated using household earnings from the SIPP. Figure (2.7) shows a kernel density plot of earnings in 2004, the first of the two years I consider, among families with 3 members in the six states with a Medicaid earnings threshold at 100% of FPL. The vertical line is the earnings threshold located at \$15,670. There is a clear peak in the density around \$15,670, with some missing mass to the right, providing evidence of earnings responses to Medicaid.

Figure (2.8) depicts earnings in the SIPP in 2005 in all six states, while Figure (2.9) depicts earnings in 2005 in the three states that moved their Medicaid earnings threshold. We should expect to see a decrease in bunching in 2005, because three states moved their earnings thresholds, unless adjustment costs prevent responses to the change. The vertical line is the earnings threshold at \$15,670, adjusted for the inflation rate, which was the location of the threshold in 2004. There is still evidence of bunching at \$15,670, with little change between 2004-2005.

2.6.2 Estimates of bunching

To estimate the counterfactual earnings distribution and calculate observed bunching, I estimate equations (2.11) and (2.13) and calculate equation (2.12) among three subgroups. First, I present results from running equation (2.11) among families of size 3 in states with a Medicaid notch at 100% of FPL. Then, I check whether there is evidence for bunching in 2004 among households of size 4 with a Medicaid earnings threshold at 100% of FPL, or among households of size 3 in states with a CHIP threshold at 100% of FPL, but no Medicaid notch. This is to confirm that observed bunching near 100% of FPL is motivated by the Medicaid notch, and not by the EITC kink or CHIP threshold.

Main results

Table (2.2) contains estimates of excess mass, the result of running equation (2.11) among families of size 3 in states with a Medicaid notch at 100% of FPL. The first two panels of Table (2.2) present results using only SIPP and CPS data, respectively, while panels (3) and (4) present results using SIPP and CPS data combined. First I will discuss the results in panel (3), using the SIPP and CPS data.

My estimates of excess mass are denoted by b in Table (2.2). In 2004, the ratio of observed mass near 100% of FPL over predicted mass estimated from equation (2.11) was 1.50, meaning that the observed mass near the Medicaid notch is 50% higher than would be expected if there were no discontinuity at 100% of FPL. In 2005, in Connecticut, where the earnings threshold increased, this number was 1.43, implying that many workers continued to earn near 100% of FPL, even after the threshold moved. On the other hand, when the Medicaid threshold falls, excess mass near 100% of FPL drops to 1.03, implying that there is no excess mass remaining near the threshold. Finally, there appears to be no evidence for bunching in 2005 near 90% of the federal poverty line when the threshold falls. Graphs of the counterfactual earnings density plotted against observed density are included in Figures (2.10)-(2.12). These graphs provide visual evidence of the patterns observed in the earnings response. In 2004, there is some evidence of excess mass near 100% of FPL, and this excess mass is still evident in 2005 when the threshold increases. In 2005 when the threshold decreases, there appears to be little or no excess mass, which is depicted in Figure (2.12).

In Table (2.2), I also present the standard error for each estimate of excess mass, or b, in parentheses, which are calculated using a bootstrapping procedure described in the estimation strategy section. Excess mass is statistically significant if it is greater than 1, as a value of 1 means that there are exactly as many people earning near the threshold as we would expect if there were no discontinuity in the budget constraint. Excess mass is statistically significant in columns (1), and is not statistically significant and imprecisely estimated in columns (2), (3), and (4).

The earnings responses that are implied by the calculated excess mass are also included in Table (2.2) and range from \$427.34 in the case of bunching at 100% of FPL when the threshold falls, up to \$802 in 2004.[8] An earnings response of \$802 in 2004 means that families earning up to \$802 above the earnings threshold at 100% of FPL are willing to reduce their earnings to qualify for benefits.

The earnings responses that I estimate are generally small. As mentioned in section (2.2), these results may be biased due to the asset test or using annual income instead of monthly income. They are also in line with similarly small or insignificant estimates from the literature, as mentioned in the introduction.

Panel (4) of Table (2.2) contains estimates of bunching in 2004 and 2005 allowing for round number bunching in the combined SIPP and CPS sample. These are estimates using equation (2.13). Estimates of bunching are generally smaller, and especially so for 2004, because families reporting \$15,000 are included in my bunching estimates in this particular case. The estimate of bunching in 2004 falls by around \$100, to account for the fact that some people report earnings of \$15,000 simply because it is a round number.

EITC and CHIP bunching

As mentioned earlier, there are two programs with thresholds near 100% of FPL that could be confounding my results, EITC and CHIP. To investigate whether these programs are affecting the earnings distribution near 100% of FPL, I study the earnings distrubutions of

[8]To calculate the earnings response, I multiply the excess mass by the bin width, \$800, and add \$400, because the center bin is located between \$15,270 and \$16,070, implying that the excess mass is coming from the right of \$16,070. The formula for calculating the implied earnings response is: $e_r = (b-1) * 800 + 400$

families that are not affected by the Medicaid notch at $15,670 and present estimates of excess mass using equations (2.11) and (2.13) for these families.

To study whether the Earned Income Tax Credit is affecting the earnings distribution, I have included, in Figure (2.13), the earnings histogram of families with 4 members in the six states with a Medicaid earnings threshold at 100% of FPL in 2004. These families have 4 members, so their Medicaid threshold is near $19,000 and the EITC creates a kink in the budget constraint near $15,000. If the EITC is confounding my estimates, there should be excess mass near $15,000 among these families, which is attributed only to the EITC and not Medicaid, because the Medicaid earnings threshold is not near $15,000 for these families. The earnings histogram in Figure (2.13) is calculated using SIPP and CPS data combined, and there appears to be limited graphical evidence of excess mass near $15,000, although there is clear evidence of round number bunching.[9]

After calculating excess mass using equation (2.13) among families of size 4 in states with a Medicaid threshold at 100% of FPL, I fail to reject the null hypothesis that there is no excess mass located near the Medicaid earnings threshold at $15,670, once I control for round number bunching. These results are presented in the first row of Table (4). While the mean earnings response is not significantly smaller than that estimated among families of size 3, it is a much noisier estimate.

I also look at families of size 3 in states with a CHIP threshold at 100% of FPL, but no Medicaid notch, to study how CHIP thresholds affect the earnings distribution. If CHIP thresholds are confounding my estimates, there should be evidence of excess mass at 100% of FPL, which can be attributed to CHIP and not to Medicaid, as these families do not face a Medicaid threshold at 100% of FPL. Figure (2.14) depicts the earnings distribution of household income in 2004 among these families. From the figure, it appears that there is some excess mass located to the left of $15,670.

To evaluate whether this is statistically significant evidence of responses to CHIP thresholds, I estimate equation (2.11). These results are presented in the second row of Table (4). Among families of size 3 in states with a CHIP threshold at 100% of FPL, but no Medicaid notch, I find again no statistically significant evidence of excess mass.

[9]To relate this to Saez (2010), who finds bunching near the first kink point in the EITC schedule, I include all families of size three, including wage earners and self-employed. The results from Saez (2010) are found among self-employed taxpayers, and the author uses IRS data. As I mentioned in the literature review, tax return data is not ideal for studying Medicaid bunching, because Medicaid eligibility is determined using a different measure of income than that used for the EITC.

Put together, these pieces of evidence suggest that the decision to locate near 100% of FPL is motivated by the loss of parental Medicaid coverage at 100% of FPL, rather than the loss of insurance for children or the kink in the EITC. This coincides with the logic that states used to motivate providing coverage to parents as well as children. Both Oregon and Tennessee argued, when making the decision to introduce more generous coverage for parents, that adding coverage for parents would increase coverage for children, as families would be more likely to adopt Medicaid if the whole family received benefits.

There are also results in the literature that support the conclusion that CHIP earnings thresholds are unlikely to generate large labor supply or earnings responses. Card and Shore-Sheppard (2004) study the Children's Health Insurance Program expansions of the late 1990's to estimate the effects of Medicaid expansions on children's health insurance coverage and find very low take-up among newly eligible children. The authors use SIPP data to study these changes, which I also use in my work.

2.6.3 Elasticity approximation

I relate the estimated earnings response to the elasticity of earnings using a "reduced-form" approximation, following Kleven and Waseem (2013). This approximation allows the elasticity to be expressed as an explicit function of observables, which is not usually possible in the case of a notch in the budget constraint. The purpose of this exercise is to show how the observed earnings response is a function of both the elasticity of earnings and the value of Medicaid, and that ignoring the role of the value of Medicaid would lead to incorrect estimates of the earnings elasticity. The approximation assumes that adjustment costs are zero. The formula for the approximation is:

$$e_R = \frac{\frac{\Delta z_a^*}{z^*}}{\frac{t^* - t}{1 - t^*}}$$

This expresses the elasticity as a function of observable parameters and the estimated earnings response, Δz_a^*. Here, t^* is a simulated marginal tax rate, depicted in Figure (2.15), that creates a kink in the budget constraint at z^*, while t is the marginal tax rate. Provided that t^* and t are close together, this simulated change in the marginal tax rate will produce an earnings response that is very close to the earnings response caused by a notch of size T. I calculate t^* by calculating the slope of the line going through the initial notch and through counterfactual earnings in the presence of the notch. Using the formula for the slope of a line, t^* is given by:

$$t^* = \frac{(1-t)z^* - (1-t)(z^* + \Delta z_a^*) + T}{\Delta z_a^*} = t + \frac{T}{\Delta z_a^*}$$

Plugging in for t^* and simplifying, the reduced-form elasticity is:

$$e_R = \frac{(\Delta z_a^*)^2(1-t)}{Tz^*} \tag{2.15}$$

To show how elasticities and the notch size are related, consider the third panel of Table (2.3), which shows calculations of the elasticity approximation in equation (2.15) for different values of Medicaid, T. In 2004, families of size three with earnings near \$15,670 faced a marginal tax rate of 16% due to the EITC kink, so $(1-t) = 0.84$. The Medicaid threshold was at \$15,670, which is z^*. From my estimates of excess mass, the earnings response to the Medicaid notch is $\Delta z_a^* = 801.61$ in 2004.

If beneficiaries value Medicaid at its full cost, which was \$2,509 on average, per adult beneficiary (Kaiser Family Foundation), their elasticity of earnings is 0.019. Previous estimates of the value of Medicaid to the beneficiary suggest that it is around 25-50% of the cost to the government (Gallen 2014). If the value of Medicaid to the beneficiary is only 25% of the cost, this elasticity increases to 0.075. The logic behind this result is that, for a fixed earnings response, the elasticity of earnings that would explain that earnings response must be higher if the value of the transfer is lower. This exercise is intended to show why estimating the value of Medicaid is an important feature of my model, as the observed earnings response to the Medicaid earnings threshold depends strongly upon both.

2.6.4 Structural elasticity estimates

After implementing the minimization problem outlined in Section (2.5.3) , I obtain an estimate of the elasticity of earnings of 0.012. The value of Medicaid, T, is estimated to be \$613.82, traditional adjustments costs, a_2, are \$258.10, and the additional cost to adjust earnings while maintaining eligibility, $a_1 - a_2$, is estimated to be \$298.97. These estimates are presented in the first panel of Table (2.3).

Observed bunching near 100% of FPL does not respond when the threshold increases, and observed bunching near 90% of FPL in 2005 is both small and noisy, so traditional adjustment costs and the additional cost to adjust earnings while maintaining eligibility are rather high. My adjustment cost estimates are also noisy, primarily because my estimates of the earnings response in Table (2.2) are noisy. From the description of identification in this model, the value of Medicaid is identified by differences in how excess mass near the initial earnings threshold changes when the threshold increases or decreases. Adjustment costs are used to explain differences in bunching in each case that cannot be attributed to the value of Medicaid. As a result, noise in estimates of excess mass contribute to noise in estimates of adjustment costs.

To place my parameter estimates in the literature, a low elasticity of earnings is not a surprising result among low-income populations. For instance, Saez (2010) finds limited responses to EITC kink points, and correspondingly small elasticities, among low-income earners.

A cash-equivalent value of Medicaid of $613.82 implies that beneficiaries value Medicaid coverage at around 24% of its cost to the government. This is similar to estimates of the value of Medicaid from the literature. Gallen (2015) uses a structural model of labor supply in the context of the Tennessee Medicaid program to estimate a value for Medicaid of about $0.26 per dollar spent on expenditures. The author uses enrollment responses to changes in the premium structure used by the state's waiver program to identify his parameters of interest. Similarly, Finkelstein, Hendren, and Luttmer (forthcoming) use evidence from the Oregon Health Insurance Experiment to estimate the value of Medicaid, and find a net welfare benefit of between $0.20 and $0.50 per dollar of government spending.

My estimate is on the low side of those found in the literature, but this seems reasonable, as the population I study is non-disabled adults with positive earnings. The costs of Medicaid are highly skewed towards disabled and aged recipients, with low-income, non-disabled adults being the least costly group to insure. It is reasonable that working adults with low medical costs would have a value for Medicaid that is on the low side among all Medicaid beneficiaries.

The second panel of Table (2.3) presents estimates of my parameters of interest using earnings responses that control for round number bunching. The value of Medicaid falls to $553.32, which is around 22% of its cost to the government. The value of Medicaid falls because the estimated earnings response to the Medicaid notch in 2004 is smaller, while the other three estimated earnings responses are similar to those estimated without round number bunching. As discussed earlier, the value of Medicaid is identified using differences in bunching at the Medicaid threshold in 2004 at 100% of FPL and in 2005 at 90% of FPL. This difference falls using round number bunching because bunching in 2004 at 100% of FPL falls, and the value of Medicaid falls to account for this.

2.6.5 Marginal value for public funds

To interpret how the value of Medicaid and earnings responses interact with the cost of the program, I calculate the marginal value for public funds (MVPF) spent on Medicaid, following the formula developed in Hendren (2015). This calculation shows how much value is returned to beneficiaries for each dollar of expenditure, which incorporates the total cost of the benefit, including the cost of earnings responses. The author defines the fiscal externality as the additional cost of the program caused by earnings responses. In my context, the earnings responses are fairly small, so the fiscal externality of the program is also small, because the effect on tax receipts is limited.

The formula for the MVPF for Medicaid is:

$$\text{MVPF} = \frac{WTP}{Cost + FE}$$

As a reminder, the average cost for adult, non-disabled Medicaid recipients in 2004 was \$2,509. Taxpayers earning near \$15,670 in 2004 faced a marginal tax rate of 16%, as they are in the phase out region of the EITC. This means that the fiscal externality to behavioral responses to Medicaid are amplified, as each additional dollar of earnings is associated both with an increase in tax receipts and a decrease in tax credits.

The willingness to pay for Medicaid that I estimate above is \$613.82, and the earnings response in 2004 was \$801.61 The fiscal externality, or FE, is the reduced tax receipts caused by earnings responses, which is $\$128.27 = 0.16 * \$801.61$, or 16% of the earnings response. This means the marginal value for public funds is $\frac{613.82}{128.27+2509} = 0.23$. The relatively small earnings response, combined with the low marginal tax rate faced by households earning near the Medicaid threshold, means that the marginal value for public funds is very close to the ratio of the value for Medicaid divided by the cost of the benefit.

Relative to the MVPF for other programs calculated by Hendren (2015), this value is low. The Food Stamps program had a MVPF of between 0.53 and 0.66, based on a reduction of tax receipts of $588 for an average $1,153 benefit and a value for Food Stamps of between $0.87 and $1 for every $1 of expenditure. These calculations are based on labor supply responses estimated by Hoynes and Schanzenbach (2011), who find an imprecise but very large earnings response to Food Stamps. The head of the household is estimated to reduce earnings by $4,418 for a benefit of up to $526 per month. The reduction in family income is associated with a reduction of tax receipts of $588, the fiscal externality. As a result, while Food Stamps are a valued at very close to their cost, they also generate large fiscal externalities, which drives down the MVPF of Food Stamps.

Section 8 housing vouchers were calculated to have a MVPF of 0.78. This calculation is based on estimates from Jacob and Ludwig (2012) and Reeder (1985). Jacob and Ludwig (2012) calculate a maximum voucher of $8,235 for their population and use a wait list lottery to estimate an earnings response of $910 per year. Reeder (1985) estimates a benefit of $83 for each $100 spent in the program. This means that the MVPF is very high, because housing vouchers are valued at near their cost, and generate small labor supply responses relative to the total cost of the program. Compared to both of these programs, the MVPF of Medicaid is low, primarily due to the low value of Medicaid relative to its cost.

2.6.6 Heterogeneity

One feature of my model is that it does not allow for heterogeneity in parameters, meaning that I assume the same elasticity of earnings, adjustment costs, and value of Medicaid for everyone. To understand how this restriction might affect my estimates, I look at how bunching and the associated earnings response in 2004 varies across two observable characteristics.

First, I look at differences in families with children under 6 and children over 6. The earnings distributions and associated counterfactual densities are presented in Figure (2.16), and the second panel of Table (2.4) presents the associated earnings response. Unsurprisingly, families with children under 6 have lower incomes, and are thus more likely to earn near the Medicaid threshold, meaning that they drive my baseline estimates. This is evident in Figure (2.16), where the earnings distribution of families with children under 6 is much higher near $15,000 than the earnings distribution of families with children over 6. The excess mass in panel (2) of Table (2.4) is also much higher, at 1.85, for families with children under 6 than it is for children over 6. The excess mass for families with children over 6 is also not statistically significant.

Second, I look at differences in the earnings response by marital status. Figure (2.17) presents the earnings distributions and associated counterfactual densities, and the third panel of Table (2.4) presents the earnings response. Again, unsurprisingly, unmarried families are the most likely to earn near the Medicaid earnings threshold. Unmarried families also respond more to the Medicaid notch, with an earnings response of \$907.70, compared to an earnings response of \$763.04 for married families.

2.7 Conclusion

In this paper, I use non-linear budget set methods to estimate earnings responses to in-kind transfers, and use these results to estimate the value of the transfer. I estimate a cash-equivalent value of Medicaid of \$613.82 and adjustment costs of \$258.10 to adjust earnings and an additional \$298.97 to adjust earnings while maintaining eligibility. This value of Medicaid and earnings response translates to a marginal value for public funds of 0.23. My estimate of the value of Medicaid is similar to those found in the literature, around 25% of the cost of provision.

A few of the limitations of my empirical strategy are that I use self-reported data in a relatively small sample, which adds noise to my estimates of the earnings response. I also do not account for some features of the Medicaid program, like asset tests and eligibility re-verification periods, which adds measurement error to my estimates, as I am not using the appropriate measure of income for Medicaid eligibility in states that implement asset tests and eligibility re-verification periods of less than one year. Some of these limitations could be overcome using administrative data, which would reduce measurement error from self-reported earnings and measurement error from using annual income. Using administrative data would not resolve the issues caused by asset tests, however. The benefit of using self-reported data is that I have access to demographic characteristics like children's ages, which are likely to affect the value of Medicaid.

Future work on this topic that I plan to pursue includes extending the model to incorporate heterogeneity in parameters and exploring the use of administrative data to see how the results differ when more precise earnings data is available. There are two strategies I could pursue to incorporate heterogeneity in parameters. First, to avoid putting structure on the distribution of parameters, I may consider a partial identification strategy, following Kleven and Waseem (2013). Alternatively, to identify all four of my key parameters, I could put some structure on the distribution of parameters. This is necessary because there is no linear relationship between the earnings response and my parameters of interest. This means that the earnings responses that I estimate cannot be expressed as a function of the expected values of my parameters of interest, and moments of the joint distribution appear in the indifference conditions, when parameters are expressed as random variables. To overcome this, I could potentially assume some form of independence between my parameters.

My estimates suggest that households are not very responsive to the Medicaid notch, and that this is in part because they have a low value for Medicaid benefits that is much smaller than its cost of provision. However, Medicaid benefits may generate externalities that overcome this apparent inefficiency. For example, Medicaid benefits may provide access to preventive care that reduces the cost of future illnesses. Medicaid beneficiaries may be more myopic than the central planner, and so it improves efficiency to offer subsidized access to preventive treatments. This example illustrates why my estimate of the value of Medicaid is only one piece of evidence towards assessing the efficiency of the Medicaid program.

2.8 Appendix

2.8.1 Sensitivity of bunching estimates

I consider the sensitivity of my bunching estimates to three parameter choices: the width of the bins used to calculate the earnings density, the number of bins excluded in estimation, and the degree of the polynomial used to calculate the counterfactual earnings density. In every case, I use the SIPP and CPS combined dataset. I present the results of this analysis in Table (2.5).

First, I estimate bunching in 2004 by excluding 3 bins in estimating equation (2.11). This means that households earning between $13,670 and $16,070 are considered as "bunching". From Figure (2.10), there is no graphical evidence suggesting that the bunching interval is this wide, and when a third bin is excluded in equation (2.11), my estimate of the earnings response falls by over $300, although it is still statistically different from counterfactual earnings.

Next, I estimate bunching in 2004 using bins of width \$1,000. In this case, the estimated earnings response is \$822.32, which is very similar to the baseline estimate. In this case, households earning between \$14,270 and \$16,170 are considered as “bunching”. Again, the excess mass is still statistically significant.

Finally, I show how my estimates of bunching are affected by altering the degree of the polynomial used to estimate counterfactual earnings. In Figures (2.10)-(2.12), I graph counterfactual earnings distributions estimated as third and fifth-order polynomials, as well as the seventh-order polynomial I use in my baseline estimates. From the graphs, it is clear that a fifth-order polynomial produces very similar results to a seventh-order polynomial, while a third-order polynomial produces a larger estimate of the earnings response in 2004.

2.9 Tables

Table 2.1: Summary statistics

	2004 SIPP	2004 CPS
All households in survey		
Number of observations	94,625	210,649
Number of households	48,201	76,439
Mean household annual income	$52,209	$57,729
Share married households	0.48	0.52
Share with kids under 6	0.16	0.18
Average age of adults	48	46
Share working adults	0.67	0.63
3 member households		
Number of observations	13,563	40,486
Number of households	8,337	13,720
Mean household annual income	$62,621	$67,049
Share married households	0.65	0.61
Share with kids under 6	0.32	0.22
Average age of adults	43	40
Share working adults	0.75	0.69
3 member households in 6 states		
Number of observations	2,357	6,854
Number of households	1,469	2,318
Mean household annual income	$63,875	$69,871
Share married households	0.64	0.60
Share with kids under 6	0.32	0.22
Average age of adults	43	41
Share working adults	0.74	0.67

The six states are: California, Connecticut, Delaware, Ohio, Oregon, and Tennessee

Table 2.2: Normalized excess mass, or b, and implied earnings response

		(1)	(2)	(3)	(4)
SIPP	b	1.58	1.49	1.18	0.48
		(0.20)	(0.47)	(0.32)	(0.56)
	Implied earnings response	\$864.38	\$793.20	\$543.98	\$-17.99
		(160.22)	(376.43)	(258.55)	(451.13)
CPS	b	1.44	2.10	1.01	1.27
		(0.17)	(0.48)	(0.24)	(0.39)
	Implied earnings response	\$748.55	\$1,279.0	\$406.91	\$614.19
		(136.32)	(384.92)	(189.35)	(311.26)
SIPP + CPS	b	1.50	1.43	1.03	1.17
		(0.15)	(0.32)	(0.22)	(0.30)
	Implied earnings response	\$801.61	\$744.32	\$427.34	\$536.42
		(120.03)	(213.97)	(178.69)	(237.82)
SIPP + CPS:RNB	b	1.37	1.40	1.01	1.22
		(0.12)	(0.23)	(0.20)	(0.33)
	Implied earnings response	\$695.57	\$722.00	\$411.20	\$576.38
		(95.17)	(153.80)	(158.64)	(296.97)

Columns:
(1) 2004 at 100% of FPL
(2) 2005 at 100% of FPL when threshold increases
(3) 2005 at 100% of FPL when threshold decreases
(4) 2005 at 90% of FPL
Notes:
(a) Standard errors in parentheses
(b) $b = \frac{\text{Observed earnings}}{\text{Counterfactual earnings}}$; calculated using equation (2.12) from estimates from equations (2.11) and (2.13)
(c) The formula for calculating the implied earnings response is: $e_r = (b-1) * 800 + 400$
(d) RNB: Round number bunching

Table 2.3: Parameter estimates

	(e)	(a_2)	(T)	$(a_1 - a_2)$
SIPP + CPS	0.012	\$258.10	\$613.82	\$298.97
	(0.002)	(152.41)	(123.95)	(242.03)
SIPP + CPS:RNB	0.010	\$276.17	\$553.32	\$303.45
	(0.003)	(143.28)	(132.57)	(206.38)
Elasticity approximation implied by $\Delta z_a^* = 801.61$				
T	\$2,509	\$1,245.50	\$836.33	\$627.25
e_R	0.019	0.036	0.056	0.075

Notes:
(a) Standard errors in parentheses
(b) e: elasticity of earnings
(c) a_2: adjustment cost
(d) T: cash-equivalent value of Medicaid
(e) $a_1 - a_2$: additional cost to adjust earnings while maintaining benefits
(f) RNB: Round number bunching
(g) Δz_a^*: earnings response from panel (3), column (1) in Table (2.2)
(h) e_R: elasticity approximation from Section (2.6.3)

Table 2.4: Estimates of excess mass, b: alternative specifications

				2004 at 100% FPL
		Baseline estimate	b	1.50 (0.15)
			Earnings response	\$801.61 (120.03)
Bunching at 100% of FPL with no notch	(1)	4 member families at \$15,670: RNB	b	1.29 (0.27)
			Earnings response	\$629.98 (145.92)
		3 member families in CHIP states	b	1.23 (0.13)
			Earnings response	\$586.22 (103.74)
Heterogeneous bunching estimates	(2)	With kids under 6	b	1.85 (0.14)
			Earnings response	\$1082.80 (115.50)
		No kids under 6	b	1.32 (0.18)
			Earnings response	\$656.34 (147.25)
	(3)	Married	b	1.45 (0.28)
			Earnings response	\$763.24 (226.03)
		Not married	b	1.63 (0.14)
			Earnings response	\$907.70 (112.52)

Notes:
(a) Standard errors in parentheses
(b) $b = \frac{\text{Observed earnings}}{\text{Counterfactual earnings}}$; calculated using equation (2.12) from estimates from equations (2.11) and (2.13)
(c) The formula for calculating the implied earnings response is: $e_r = (b-1) * 800 + 400$
(d) RNB: Round number bunching

Table 2.5: Estimates of excess mass, b: Sensitivity analysis

		2004 at 100% FPL
Baseline estimate	b	1.50
		(0.15)
	Earnings response	\$801.61
		(120.03)
3 bins excluded	b	1.25
		(0.12)
	Earnings response	\$596.54
		(96.69)
2 bins excluded; bin width \$700	b	1.56
		(0.18)
	Earnings response	\$851.35
		(144.12)
2 bins excluded; bin width \$1,000	b	1.53
		(0.13)
	Earnings response	\$822.32
		(106.44)
2 bins excluded; bin width \$1,200	b	1.48
		(0.21)
	Earnings response	\$780.41
		(167.53)
Polynomial of degree 5	Earnings response	\$797.38
		(108.22)

Notes:
(a) Standard errors in parentheses
(b) $b = \frac{\text{Observed earnings}}{\text{Counterfactual earnings}}$; calculated using equation (2.12) from estimates from equations (2.11) and (2.13)
(c) The formula for calculating the implied earnings response is: $e_r = (b-1) * 800 + 400$

2.10 Figures

Figure 2.1: Transfer program earnings threshold

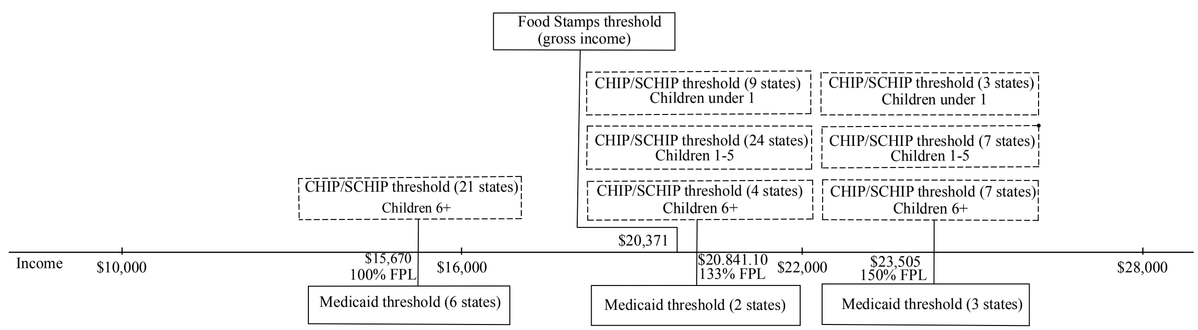

Notes

1. 3 states have Medicaid thresholds between 100% FPL and 150% FPL. California, Delaware, and Connecticut have earnings thresholds at 107% (CA, CT) and 117% (DE) for working applicants
2. Illinois has a Medicaid threshold at 140% FPL for working applicants
3. South Dakota and Arizona set at least one CHIP/SCHIP threshold at 140% FPL.
4. 33 states have Medicaid thresholds below 100% FPL.

Figure 2.2: Medicaid coverage

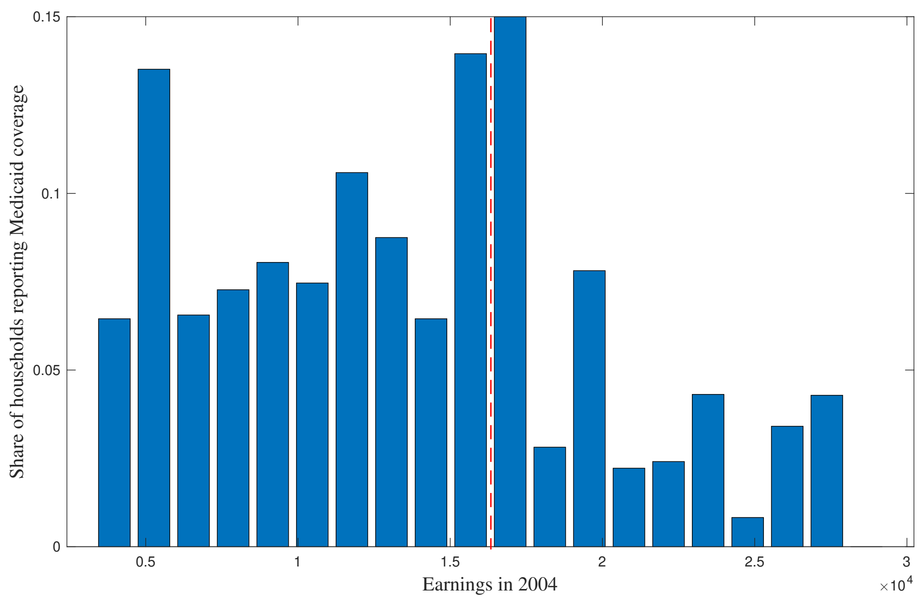

Notes:
(a) SIPP data
(b) Vertical line: 100% of FPL
(c) Families of size 3

Figure 2.3: Budget constraint

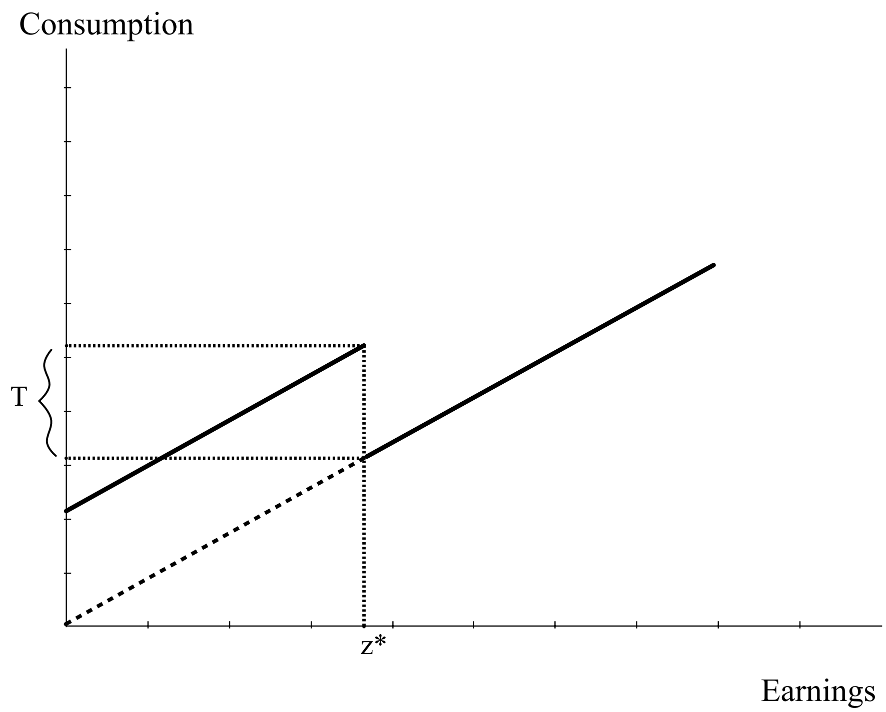

(a) Notch at z^*

Figure 2.4: Indifferent Individual

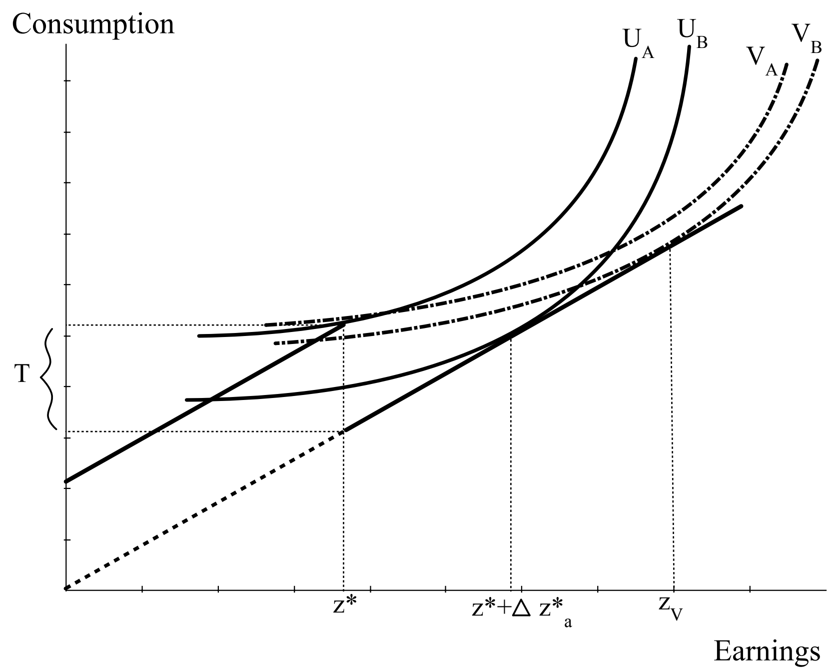

Notes:
(a) Indifferent individual U has counterfactual earnings $z^* + \Delta z_a^*$
(b) Individual V not willing to move to z^* with adjustments costs a

Figure 2.5: Two period bunching: earnings threshold increases

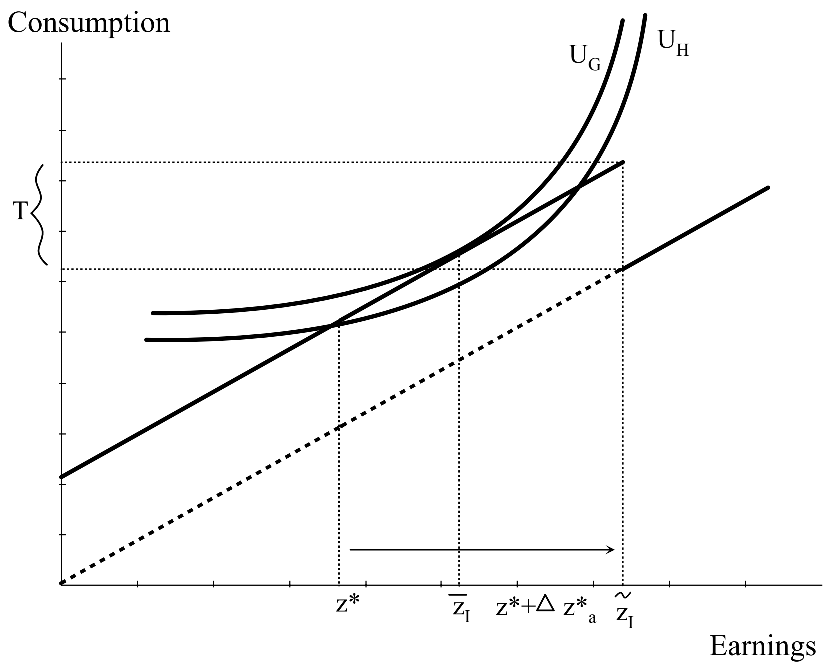

Notes:
(a) Indifferent individual U has counterfactual earnings $\bar{z}^I$
(b) Bunching interval $(z^*, \bar{z}^I)$

Figure 2.6: Two period bunching: earnings threshold decreases

(a) Lower bound

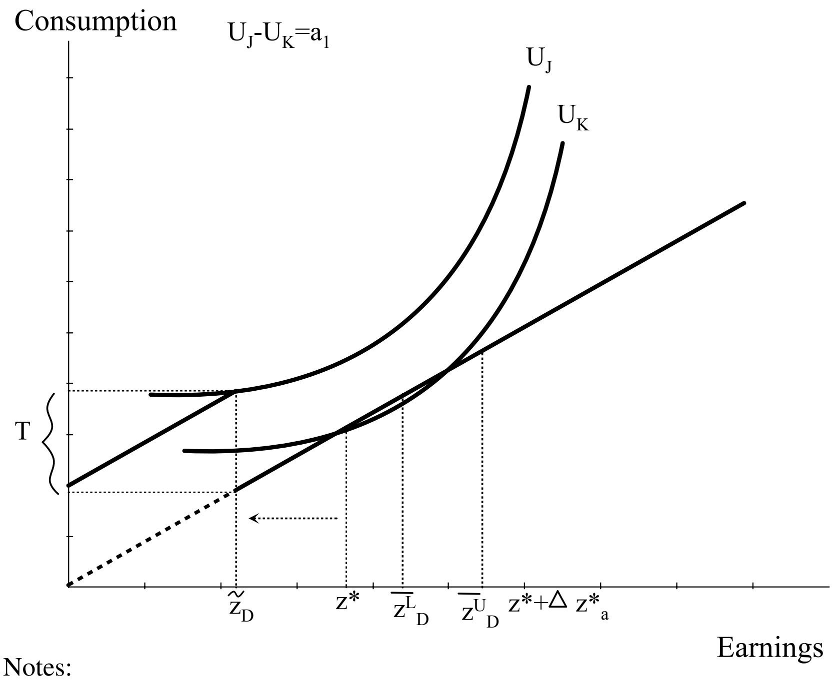

Notes:
(a) U_J: Indifference curve through new notch
(b) U_K: Indifference curve through original notch
(c) Indifferent individual U has counterfactual earnings $\bar{z}_D^L$

(b) Upper bound

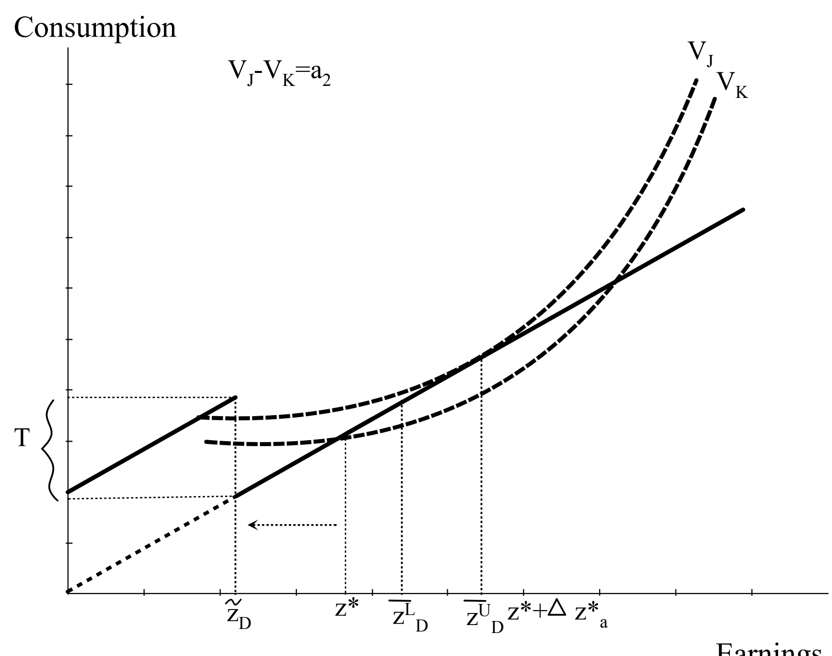

Notes:
(a) V_J: Indifference curve through counterfactual earnings
(b) V_K: Indifference curve through original notch
(c) Indifferent individual V has counterfactual earnings $\bar{z}_D^U$

Figure 2.7: Earnings in 2004

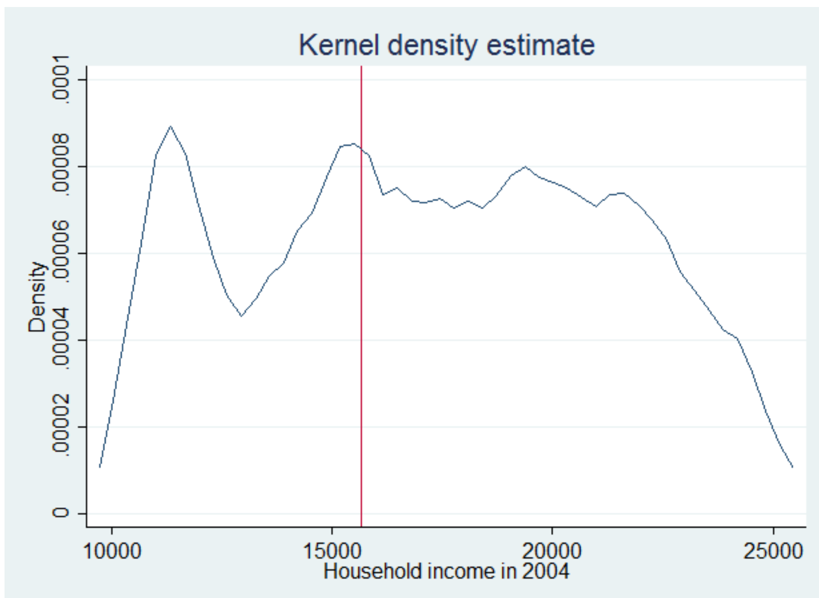

Notes:
(a) SIPP data
(b) Vertical line: 100% of FPL
(c) Excess mass: to left of vertical line

Figure 2.8: Earnings in 2005

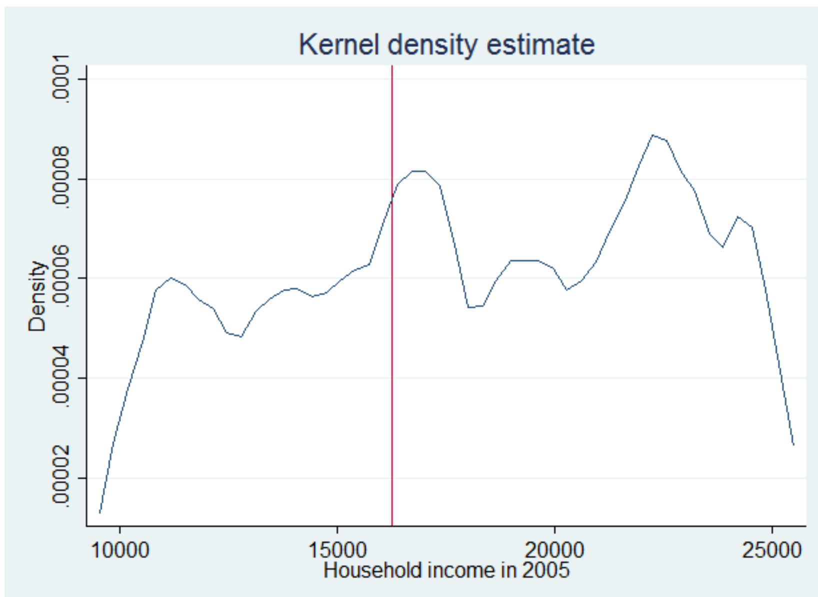

Notes:
(a) SIPP data
(b) Vertical line: 100% of FPL
(c) Excess mass: to left of vertical line

Figure 2.9: Earnings in 2005:States that moved threshold

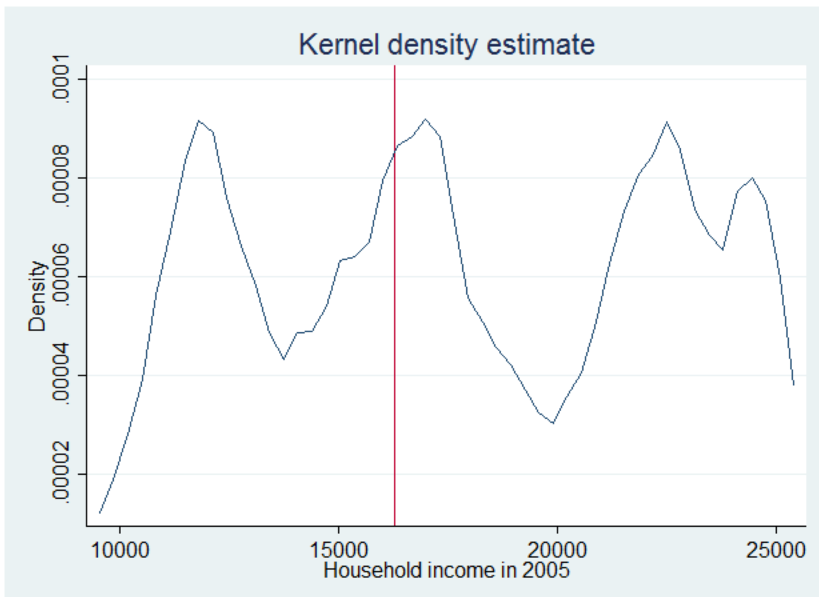

Notes:
(a) SIPP data
(b) Vertical line: 100% of FPL
(c) Excess mass: to left of vertical line

Figure 2.10: Counterfactual earnings in 2004

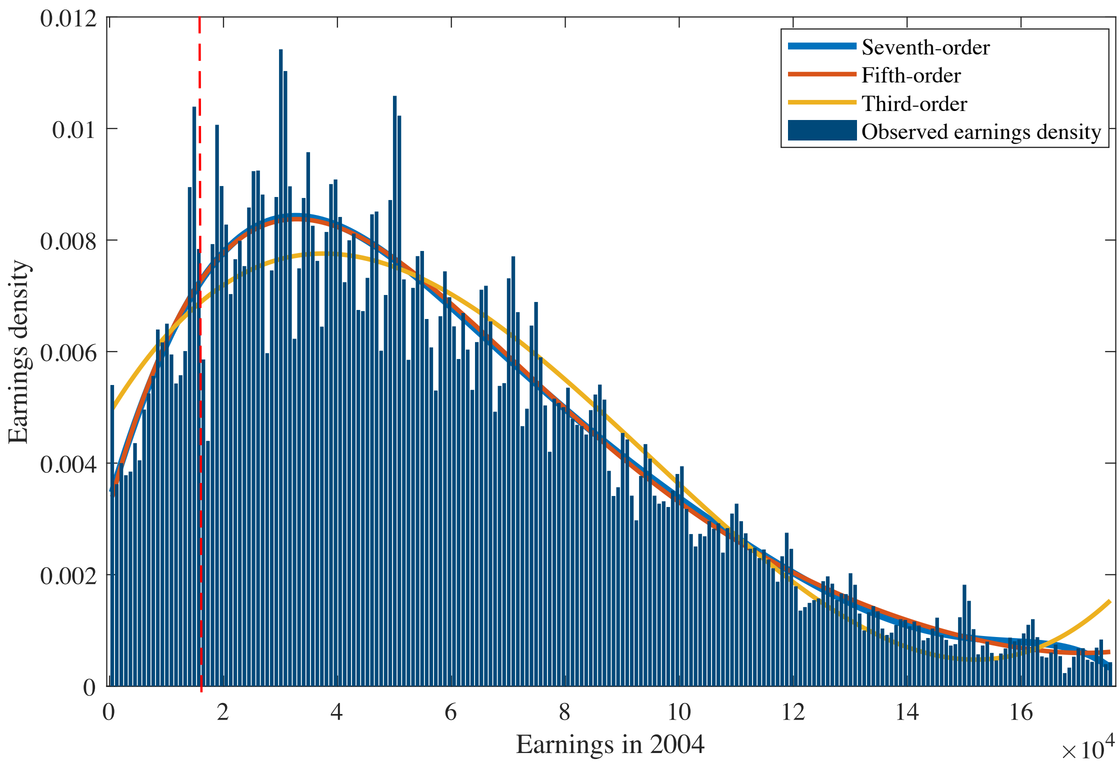

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.11: Counterfactual earnings in 2005 when threshold increases

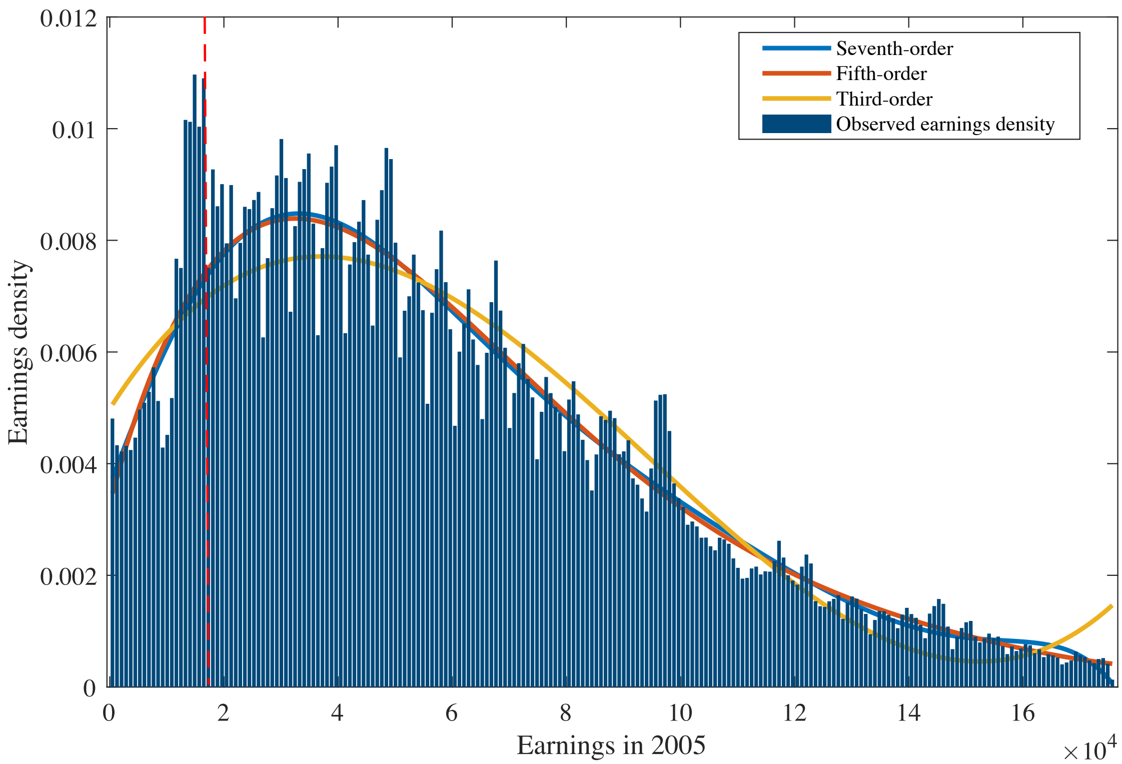

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.12: Counterfactual earnings in 2005 when threshold decreases

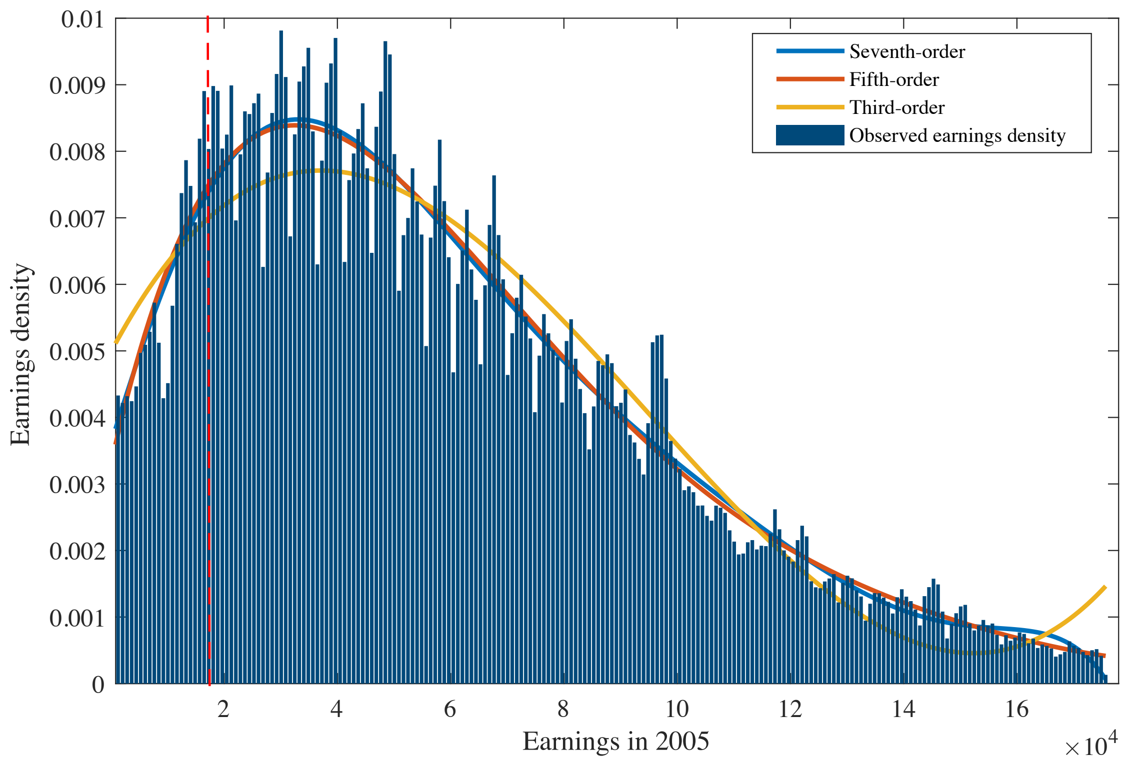

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.13: Earnings in 2004 among families of size 4

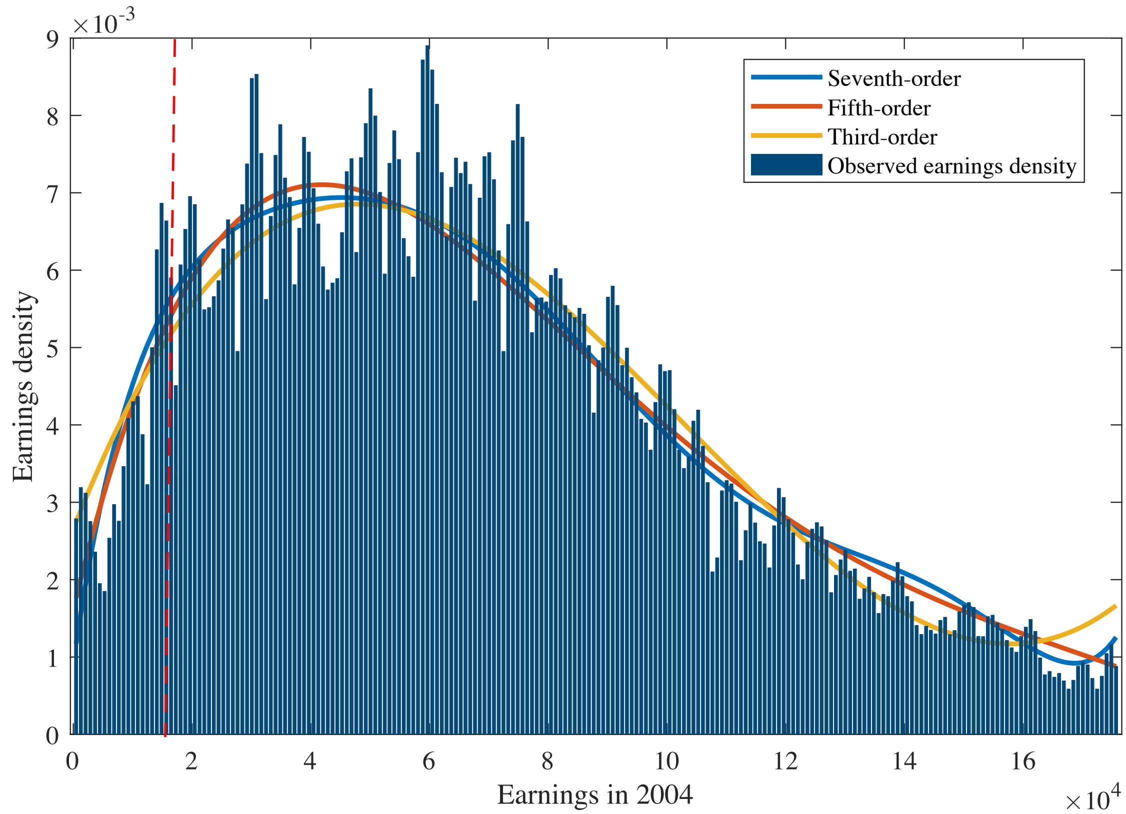

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.14: Earnings in states with CHIP threshold at 100% FPL

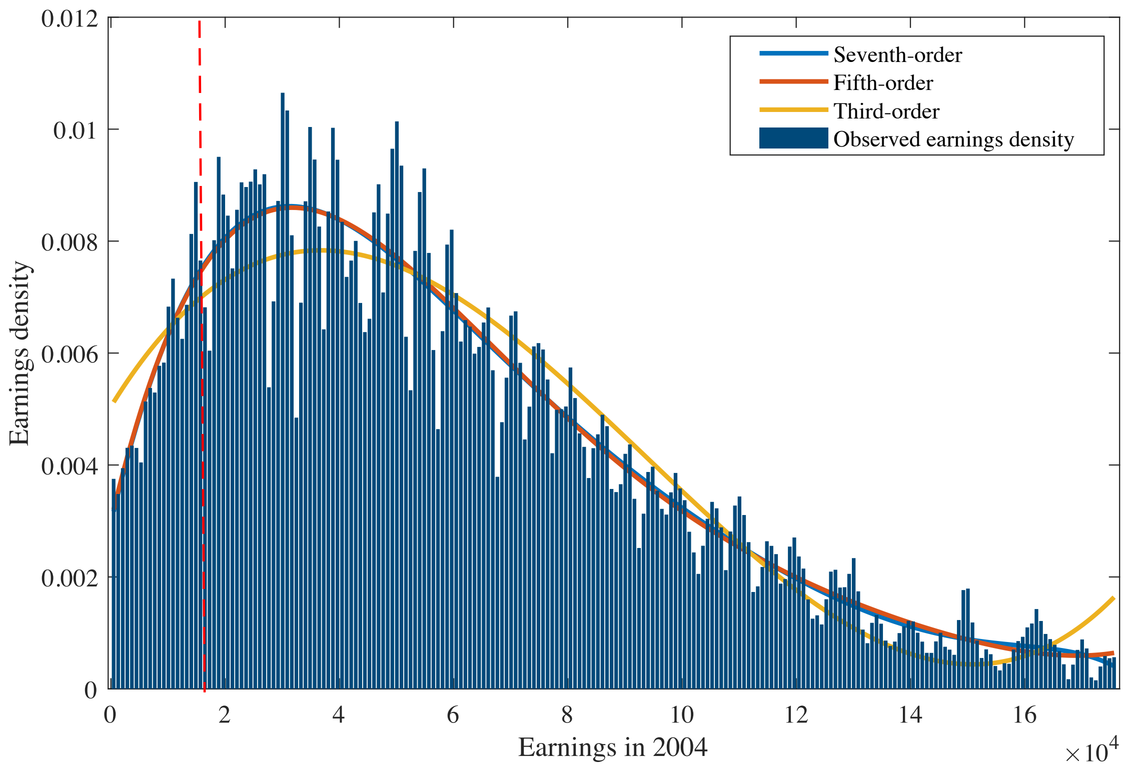

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.15: Budget constraint with simulated marginal tax change, t^*

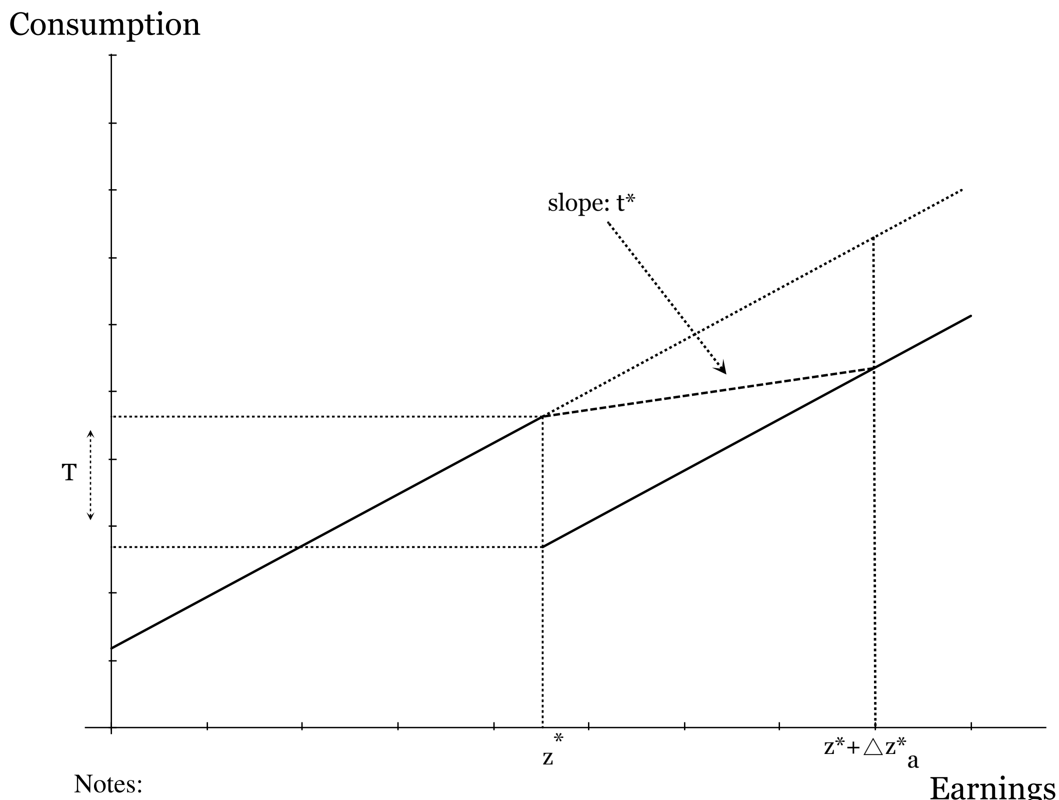

Notes:
(a) T: size of notch
(b) Δz_a^*: earnings response with notch of size T
(c) t^*: simulated marginal tax change that generates earnings response Δz_a^*

Figure 2.16: Earnings distribution by children's age

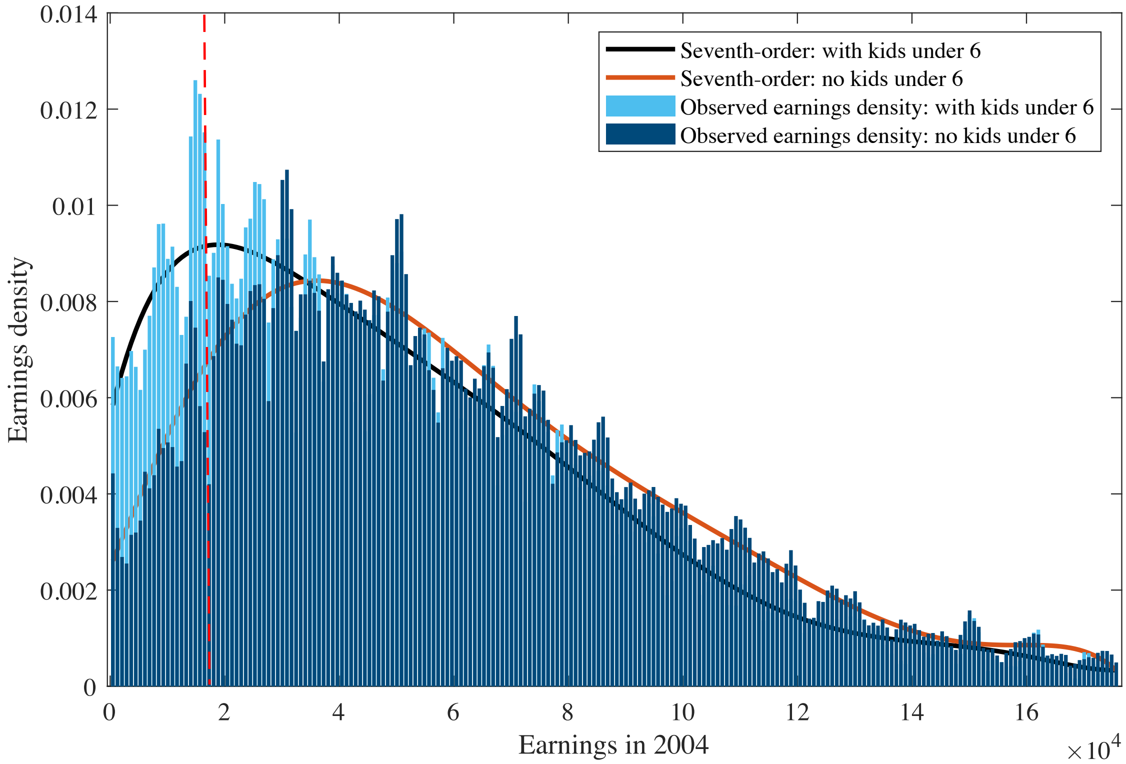

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.17: Earnings distribution by marital status

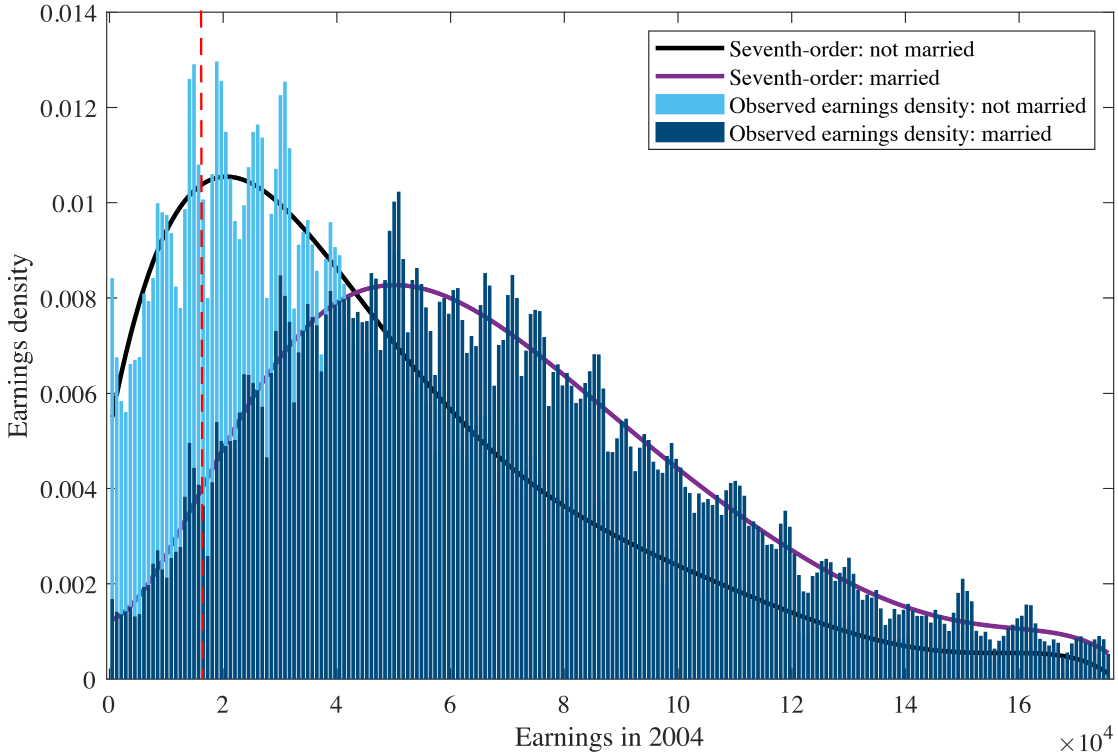

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Figure 2.18: All other states in 2004

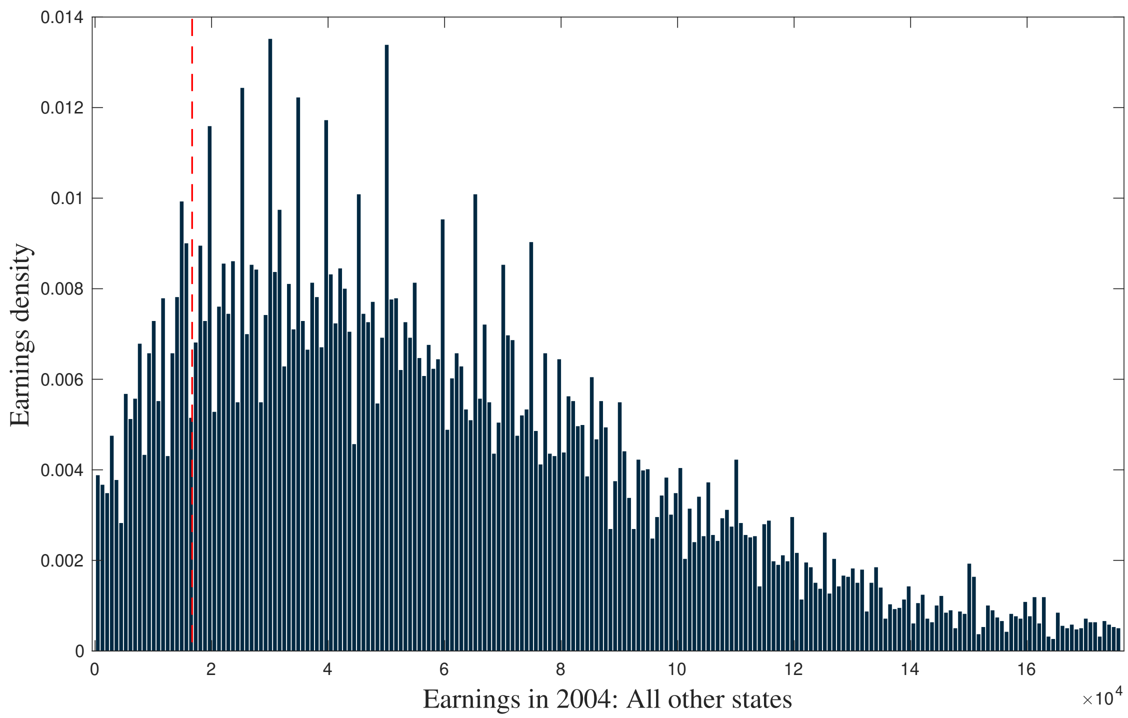

Notes:
(a) Combined SIPP and CPS data
(b) Earnings density bins smoothed using Savitzky-Golay filter
(c) Vertical line: 100% of FPL
(d) Excess mass: to left of vertical line
(e) Missing mass: to right of vertical line

Chapter 3

Assessing the Effect of Firearms Regulations Using Partial Identification Methods: A Case Study of the Impact of Stand Your Ground Laws on Violent Crime (with John Pepper)

3.1 Introduction

Despite a large literature in the social sciences and public health, empirical research has struggled to reach consensus about the impact of firearms regulations on crime.[1] Consider, for example, the recent research on Stand Your Ground (SYG) laws that allow a person to use lethal force in self-defense in places outside of the home without first attempting to retreat. Using repeated cross sectional data on annual state crime rates, recent studies have examined the impact of these laws on murder and other violent crimes. Unfortunately, this research has been inconclusive, with some studies finding positive effects, others reporting negligible/insignificant estimates, and still others concluding that SYG laws decrease

[1] See, e.g., National Research Council (NRC), *Firearms and Violence: A Critical Review,* ed., Wellford, C., J. Pepper, and C. Petrie, National Academy Press, Washington, D.C., (2005); Charles F. Manski & John V. Pepper, *How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions*, Review of Economics and Statistics, 100.2: 232-244, (2018).

violent crime.[2] Lott, for example, concludes Stand Your Ground laws reduce murder rates by 9 percent and overall violent crime by 11 percent, while Cheng and Hoekstra find that these laws increase the murder rate by 8 percent.[3] As in many other areas of research on the impact of gun regulations, empirical results on SYG laws are highly variable and sensitive to minor variations in the data and/or model.

The fundamental difficulty in drawing inferences on the effects of gun regulations is that the outcomes of counterfactual policies are unobservable. Data alone cannot reveal what the murder rate in a state with a SYG law would have been had the state not adopted the statute. To address this selection problem, observed crime data must be combined with assumptions to enable inference on counterfactual outcomes. Yet, the assumptions needed to identify these counterfactual outcomes cannot be tested empirically and different assumptions can yield different inferences.

With this in mind, we seek to make transparent how assumptions shape inferences on the effects of firearms regulations on violent crime. To do this, we apply the partial identification approach developed by Manski and Pepper to re-examine the empirical analysis of the average effect of SYG laws on violent crime.[4] These weaker, more credible models bound the average effect of SYG laws on violent crime, where the width of estimated bounds reflects the uncertainly resulting from the selection problem.

To focus attention on the sensitivity of inferences to the underlying identifying assumptions, we make two simplifying restrictions. First, we examine only the effects of adopting SYG laws in a single year rather than at any point in time. In particular, to simplify the analysis, we draw inferences on the effect of SYG laws on average violent crime rates in 2008-2010 for the thirteen states that adopted these statutes in 2006. By focusing on the impact of adopting a SYG law in 2006, we do not need to make assumptions about how the effect of the statute varies with time.[5] Second, we do not provide measures of

[2] Theory provides little guidance. These laws might deter some crimes if potential offenders perceive the costs of committing crimes may be higher. Yet, these laws may increase the lethality of criminal encounters.

[3] See John R. Lott, *More Guns, Less Crime: Understanding Crime and Gun-Control Laws*. Chicago: University of Chicago Press, (2010); Cheng Cheng & Mark Hoekstra, *Does strengthening self-defense law deter crime or escalate violence? Evidence from expansions to castle doctrine*. Journal of Human Resources, 48(3), 821-854, (2013). Also see Chandler B. McClellan & Erdal Tekin, *Stand your ground laws, homicides, and injuries. Journal of Human Resources, 52(3), 621-653, (2017).*

[4] See Manski & Pepper, supra note 1; Charles F. Manski & John V. Pepper, *Deterrence and the Death Penalty: Partial Identification Analysis Using Repeated Cross Sections*, Journal of Quantitative Criminology, 29(1), 123-141, (2013).

[5] The effects of right-to-carry laws on crime, for example, have been found to vary over time. See supra

statistical precision (e.g., standard errors or confidence intervals).[6] Instead, we view the states as the population of interest, rather than as realizations from some sampling process. Thus, imprecision expressed through the width of the bounds only reflects the selection problem, not sampling variability. We do this to focus attention on the selection problem. However, even if we wanted to provide measures reflecting the uncertainty generated by sampling variability, it is not clear how to do so in this setting where the state crime rate is the outcome variable of interest. The conventional assumption of random sampling from an infinite population is not natural when considering states as units of observation, and it is not clear what type of sampling process would be reasonable to assume.[7] [8]

We begin in Section 3.2 by demonstrating the sensitivity of inferences to the traditional models which have been used to point identify the average treatment effect (ATE). The traditional approach for resolving the selection problem is to impose assumptions strong enough to yield a definitive finding. In this case, an average effect is said to be *point-identified*. The problem is that these strong assumptions may be inaccurate, yielding flawed and conflicting conclusions.[9] We have seen this repeatedly in the empirical literature on gun regulations, including the analyses of SYG laws. In particular, the traditional assumptions asserting that expected counterfactual outcomes are invariant across geography or time lead to qualitatively different empirical findings.

In light of these conflicting findings, what credible conclusions about the effect of SYG statutes can be drawn from the empirical literature? At one extreme, some might take the

---

note 1. Also see John J. Donohue et al., *Right-to-Carry Laws and Violent Crime: A Comprehensive Assessment Using Panel Data and a State-Level Synthetic Control Analysis*, Journal of Empirical Legal Studies, vol 16(2), pages 198-247, (2019).

[6]See Manski & Pepper, supra note 1, at pp.234-235.

[7]This point is made in Manski and Pepper, id. Also see Alberto Abadie et al.,*Finite Population Casual Standard Errors*, NBER working paper no. 20325, (2014). In the setting of a randomized experimental design, they develop an alternative conceptualization for drawing inferences when one observes the entire population. However, their modified approach is not applicable in the more general observational data settings where the treatment (e.g., RTC laws) may be endogenous

[8]In the panel data literature using annual state or county crime rates to infer the impact of firearms regulations on crime, some researchers report standard errors that allow for arbitrary correlation within at a state or county – so called state/county clustered standard errors – while others do not allow for such correlations. The NRC, supra note 1 and Alberto Abadie et al., *When should you adjust standard errors for clustering?*, NBER Working Paper no. 24003, (2017), show that these clustered sampling standard errors are generally inappropriate in the standard linear panel data models (with fixed state and/or county effects) used in the literature. Manski & Pepper, supra note 1, argue that random sampling assumption underlying these conventional approaches may not be justified.

[9]See Manski & Pepper, supra note 4.

variability as evidence that empirical results are uninformative. At another, some might argue in favor of a particular model, and draw conclusions based on that model. Yet, the model assumptions cannot generally be tested empirically and often have little foundation.[10] If the assumptions lack credibility and consensus, picking a "winner" does not resolve the ambiguity.

In Section 3.3, we apply an alternative middle ground approach for drawing inferences under the weaker bounded variation assumptions developed in Manski and Pepper.[11] The basic idea is to replace the assumptions of invariance across geography or time with weaker bounded variation assumptions. For example, rather than assuming that states with and without SYG statutes would otherwise have identical crime rates, we might instead assume these two groups of states are similar to one another. Likewise, one might consider the assumption that, in the absence of a SYG statute, these two groups of states would have experienced similar but not identical trends in crime rates.

Bounded variation assumptions provide an intuitive and straightforward way to relax the traditional invariance assumptions and improve the credibility of the empirical research on the impact of gun regulations. Manski and Pepper[12] show that these assumptions partially identify the average treatment effect, yielding bounds rather than point estimates.[13]The basic insight is empirical results need not be an all or nothing undertaking. Available data and credible assumptions may lead to partial conclusions.

In Section 3.4, we draw conclusions. Under very weak assumptions, the data cannot reveal whether SYG laws increase or decrease crime. However, under our preferred set of bounded variation assumptions, we draw substantive conclusions about the qualitative

[10]See NRC, supra note 1.

[11]See Manski & Pepper, supra note 1.

[12]Id.

[13]Partial identification analysis of treatment effects from observational data was initiated in Charles F. Manski, *Nonparametric Bounds on Treatment Effects*, American Economic Review Papers and Proceedings, 80, 319–323, (1990). Charles F. Manski, *Partial Identification of Probability Distributions*, New York: Springer-Verlag, (2003) and Charles F. Manski, *Identification for Prediction and Decision*, Cambridge: Harvard University Press, 2007) provides a textbook exposition. Applications include Charles F. Manski and Daniel S. Nagin, *Bounding Disagreements About Treatment Effects: A Case Study of Sentencing and Recidivism*, Sociological Methodology, 28, 99-137, (1998); Charles F. Manski and John V. Pepper, *Monotone Instrumental Variables: with an Application to the Returns to Schooling*, Econometrica, 68, 997–1010, (2000); John V. Pepper, *The Intergenerational Transmission of Welfare Receipt: A Nonparametric Bounds Analysis*, The Review of Economics and Statistics, 82(3), 472-88, (2000); Brent Kreider et al., *Identifying the Effects of Food Stamps on the Nutritional Health of Children when Program Participation is Misreported*, Journal of the American Statistical Association, 107:499, 958-975, (2012). Also see Manski & Pepper, supra note 4.

and quantitative impact of SYG laws on violent crime. In particular, we find SYG laws have modest positive effects on rates of violent crime and murder, and uncertain effects on robbery and assault.

3.2 Inferences under point identified models

In this section, we demonstrate the sensitivity of inferences on the effect of SYG laws on violent crime, focusing on the types of point-identified models used in the literature. After providing a brief description of the data in Section 3.2.1, in Section 3.2.2 we examine the unconditional crime rates and in Section 3.2.3 we estimate the types of linear regression models applied in the literature.[14]

We consider the problem of drawing inferences on the effect of SYG laws on average violent crime rates in 2008-2010 for the thirteen states that adopted these statutes in 2006. The thirteen states are Alabama, Alaska, Arizona, Georgia, Indiana, Kansas, Kentucky, Louisiana, Michigan, Mississippi, Oklahoma, South Carolina, and South Dakota. The data reveal the annual crime rates from 2008-2010 in the thirteen states that adopted SYG statutes in 2006, but do not reveal the counterfactual crime rate that would have occurred had the thirteen treated states not adopted SYG laws. To address the selection problem, we consider three different invariance assumptions that point identify the counterfactual crime rate. In particular, we apply the assumptions that the counterfactual rate equals i) the average crime rate in the 13 treated states before 2006 (in 2000-2004); ii) the average crime rate in untreated states in the contemporary post 2006 period (in 2008-2010); and iii) the pre-treatment rate in treated states adjusted by the trend in crime rates in the untreated states, as would be the case in a difference-in-difference model.

We conclude Section 2 by examining the pre-2006 data when the thirteen treated states did not have SYG statutes. Although the identifying assumptions cannot be empirically tested in the post-2006 period when the crime rates that would have occurred without SYG laws are counterfactual, they can be assessed in the pre-2006 period. All three invariance models do not hold in the pre-treatment period.

[14] See supra note 3.

3.2.1 Data

The crime data come from the FBI's Uniform Crime Reports.[15] We use state level data on annual crime rates (per 100,000 residents) from 1977 to 2010. For each state and year, we observe the overall violent crime rate and crime rates separately for murder, assault, and robbery. We also observe whether a SYG statute is in place.[16] Our analysis uses the SYG adoption date data from Cheng and Hoekstra.[17]

To illustrate the basic counterfactual outcomes problem, we focus on the impact of SYG laws in the thirteen states that adopted this statute in 2006. We refer to the states adopting SYG statutes in 2006 as "treated" states, and those not adopting as "untreated" states. The eight other states that adopted SYG laws between 2005 and 2009 are excluded from the analysis. These states are Florida (2005), Missouri (2007), Montana (2009), North Dakota (2007), Ohio (2008), Tennessee (2007), Texas (2007), and West Virginia (2008).

Figures 1A and 1B display the annual time series of murder and robbery rates in treated and untreated states over the period 1990–2010. The figures reveal several interesting characteristics of the crime rates.

First, murder rates in the treated states exceed the analogous rates in the untreated states, while the opposite pattern exists for robbery. Second, the annual time series variation from 1990 to 2010 in crime rates for the two groups of states is similar but not identical. For example, crime rates in untreated states have a more pronounced drop during the 1990s than those in treated states.

---

[15] These panel data were originally assembled by John Lott and have subsequently been modified, corrected, and updated several times. Our analysis uses the iteration assembled and evaluated by Abhay Aneja et al., *The Impact of Right-to-Carry Laws and the NRC Report: Lessons for the Empirical Evaluation of Law and Policy*, American Law and Economic Review, 13(2), 565-632, (2011).

[16] There are significant differences in the definition of a stand your ground law within the literature. In particular, Lott, supra note 3, studies primarily "Castle Doctrine" laws, which remove the "duty to retreat" from the home before using lethal force. Once these laws are passed, people are allowed to use force to defend their homes even if they could retreat from the home to safety. More recent papers in the literature consider a Stand Your Ground law to be a law that removes the duty to retreat from somewhere outside of the home. McClellan & Tekin, supra note 3, consider any law that removes the duty to retreat from a public space, while Cheng & Hoekstra, supra note 3, consider laws that remove the duty to retreat from someplace outside of the home, which might include private property such as a vehicle. We use the classification reported in Cheng & Hoekstra, supra note 3. Thus, the thirteen states that passed a law in 2006 expanded the "Castle Doctrine" to someplace outside of the home.

[17] See Cheng & Hoekstra, supra note 3.

3.2.2 Unconditional analysis

Consider the problem of using these data to draw inference on the impact of adoption of an SYG law on the average murder rate in 2008-2010 for the thirteen states that adopted these statutes in 2006. Table 3.1 displays the average murder rate per 100,000 residents from 2000-2004 and from 2008-2010 in the states that did and did not adopt SYG statutes in 2006.

What do these data reveal about the effect of SYG laws on the 2008-2010 murder rate in treated states? The average murder rate in 2008-2010 in the thirteen states that adopted SYG statute in 2006 equals 6.17. However, the data do not reveal the murder rate that would have been observed in the treated states had the statute not been adopted in 2006. The conventional practice has been to apply assumptions that are strong enough to learn this counterfactual crime rate and, hence, the average treatment effect.

In particular, these data along with the following three simple assumptions provide three different estimates the counterfactual average murder rate:

1. *A "before-after" assumption* uses the pre-period rates among the treated states, 6.94;

2. *A "contemporaneous" assumption* uses the contemporaneous rate in the untreated states, 4.52; and

3. *A difference-in-difference assumption (DnD)* adjust the pre-period rate in the treated states by the time-series variation in the untreated state, 6.29 (=6.94-(5.16-4.52)).

In turn, these data and assumptions result in estimated treatment effects of -0.77, 1.66, and -0.12, respectively. So, depending on the model, the estimates imply that the SYG law decreases, increases or has almost no effect on murder rates.

Table 3.2 summarizes these results, and displays the estimated average treatment effect for the thirteen treated states. Estimates of the average treatment effect of SYG laws on violent crime, murder, robbery and assault are displayed in Table 3.3. For all four crime rates, the results are highly sensitive to the different modeling assumptions.

Given certain assumptions, each estimate measures the effect of SYG laws on the crime rate. The *before-after estimate* is correct under the assumption that, except for the enactment of a SYG law, no determinant of criminal behavior changed in treated states between the early and late 2000s. The *contemporaneous comparison estimate* of the 2008-2010 crime rates is correct under the assumption that, except for the presence of the SYG statutes in the treated states, the populations of the treated and untreated states had the same propensities for criminal behavior and faced the same environments. The *difference-in-difference (DnD) estimate* is correct under the assumption that, in the absence of a SYG statute, the treated and untreated states would have experienced the same change in crime rates between the early and late 2000s. Thus, each of the three estimates can be justified by specific invariance assumptions. However, it may be that none of the assumptions hold and the variation in empirical findings shows that these invariance assumptions cannot jointly hold.

Importantly, these are more than contrived illustrations of the sensitivity of inferences to different assumptions. All three of these research designs have been applied in the literature on the impact of SYG statutes, especially the before-after and DnD models.

While commonly used, these invariance assumptions may not credibly address the selection problem. SYG statutes are not likely to be randomly assigned, as would be the case in a randomized control trial, and any imaginable comparison group is likely to differ in ways that may lead to spurious correlations in the observed data and biased inferences on the impact of firearms regulations.[18] [19]

To address this concern, many of these studies (especially those using DnD models) use multiple regression models to statistically account for observed factors such as state demographics, socioeconomics, policing, other firearms laws, and so forth. In this case, researchers assume that an invariance assumption applies conditional on the set of observed covariates even if it may not apply when excluding such control variables from the analysis. Yet, the fact that states and/or time-periods with the same observed covariates have different

[18]The before-after invariance model is applied in Mitchell B. Chamlin, *An assessment of the intended and unintended consequences of Arizona's Self-Defense, Home Protection Act*, Journal of Crime and Justice, 37(3), 327-338, (2014); Mitchell B. Chamlin & Andrea E. Krajewski, *Use of force and home safety: An impact assessment of Oklahoma's Stand Your Ground Law*, Deviant Behavior, 37(3), 237-245, (2015); Ling Ren et al., *The deterrent effect of the castle doctrine law on burglary in Texas: a tale of outcomes in Houston and Dallas*, Crime & Delinquency, 61(8), 1127-1151, (2015); and David K. Humphreys, *Evaluating the impact of Florida's "stand your ground" self-defense law on homicide and suicide by firearm: an interrupted time series study*, JAMA internal medicine, 177(1), 44-50, (2017).

[19]DnD invariance models are applied in a number of papers. See supra note 3. Also see Daniel Webster et al., *Effects of the repeal of Missouri's handgun purchaser licensing law on homicides*, Journal of Urban Health, 91(2), 293-302, (2014); Mark Guis, *The relationship between stand-your-ground laws and crime: a state-level analysis*, The Social Science Journal, 53(3), 329-338, (2016); Vincent Ferraro & Saran Ghatak,*Expanding the Castle: Explaining Stand Your Ground Legislation in American States*, 2005–2012. Sociological Perspectives, 1-22, (2019).

firearms regulations suggests that confounding unobserved factors may play a role in the selection process even after controlling for observed variables. We consider the conditional analysis approach below, in Section 3.2.3.

3.2.3 Conditional analysis

Some evaluations of SYG laws analyze crime data across many states and years, and statistically control for a large set of observed covariates.[20]Although these studies use additional data on SYG statutes, annual crime rates, and covariates, the empirical findings rest on similar if not stronger identifying assumptions. Data alone cannot resolve the selection problem.

Consider, for example, the linear panel data models with state-year fixed effects. This basic panel data model relies on the strong invariance assumption that, after controlling for a set of covariates, SYG laws have the same effect on crime rates in all states and years. When combined with certain other assumptions, this homogeneity assumption has substantial identifying power. However, there is little support for the notion that the effects of firearms regulations are identical across states and time. For example, the empirical literature on right-to-carry (RTC) laws finds the effects of these laws vary over time and across states.[21]

Using repeated cross-sectional crime data from 1977-2010, we estimate a set of linear panel data models that differ in the underlying invariance restrictions, the time period used in the evaluation, the set of covariates, and the classification of SYG statutes. For each model, we present estimates for the log crime rate and the level crime rate.[22] Much of the literature studies the natural log of the crime rate but, following Manski and Pepper, our analysis using the bounded variation assumption focuses directly on the state-year crime rates. We also present results using two different sets of covariates applied in the literature on RTC laws. First, we use the control variables from Lott including detailed state race/age demographics, the arrest rate, the incarceration rate, median per-capita income, the poverty and unemployment rate, a RTC indicator, and state and year fixed effects.[23] Second, we

[20]For example, see supra note 3.
[21]See supra note 5.
[22]See Manski & Pepper, supra note 1.
[23]See Lott, supra note 3.

use the more parsimonious set of controls from Donohue et al (2019).[24] In particular, this specification does not include measures of the arrest rate or the full set of age-race demographic control variables, but does include variables measuring the size of the police force.

As with the unconditional analysis in Section 3.2.2, the estimates from these linear panel data models vary with the particular form of the invariance assumption. Table 3.4 presents results restricting the treatment group to the 13 states that adopted SYG statutes in 2006 and the control group to states that did not adopt a SYG statute. The eight states that adopted SYG laws in other years are dropped. Most but not all of the point estimates imply that SYG laws increase violent crime, but the estimated magnitudes are sensitive the underlying identifying assumption. The first two rows present basic linear difference-in-differences results with different sets of covariates. The estimates displayed in Row 1, which use the Lott specification, suggests SYG laws increased the average murder rate by 9%, the robbery rate by 2%, and overall violent crime rate by 1%, but decreased the average assault rate by 1%.[25] The second row, which reports results using the Donohue et al. specifications, suggests substantially larger effects of SYG laws; SYG laws increased murder rates by 12%, and robbery, assault and violent crime rates by about 5%.[26]

The final two rows in Table 3.4 present estimates for a before-after and contemporaneous identification strategy for the effect of SYG laws on crime rates. The before-after comparison uses only states that passed the law in 2006 and compares the murder rates in these states before and after 2006, while the contemporaneous estimate compares murder rates after 2006 in states that passed a SYG law with the murder rates in states that did not. While these estimates generally imply SYG laws increase violent crime, the magnitudes vary across the different models. For example, SYG laws are estimated to increase the robbery rate by 21% under the contemporaneous invariance model and 9% under the before-after model.

Overall, these results suggest that SYG laws increase violent crime rates for the thirteen treated states in 2008-2010. However, there is substantial variation in the magnitude of the estimated effects. Depending on the identifying assumptions, the log-crime rate regres-

[24] See Donohue et al., supra note 5. This is similar to the set of controls used in Cheng & Hoekstra, supra note 3.

[25] See Lott, supra note 3.

[26] See Donohue et al., supra note 5.

sion estimates imply that SYG laws increase murder rates between 6% and 12%, robbery between 2% and 21%, the assault rate between -1% and 17%, and the violent crime rate between 1% and 15%.

Table 3.5 expands the treatment group to all states that adopted SYG statutes between 2005 and 2010. This replicates the basic models and time period examined by Cheng and Hoekstra.[27] As in Cheng and Hoeskstra, the resulting estimates imply that SYG laws increase some crimes but decrease others. In particular, when using the Donohue et al. covariates[28] , SYG laws are estimated to increase the murder rates by 4% but decrease the robbery and assault rates by 6% and 2%, respectively.

Finally, Table 3.6 uses data from 1977-2005 to examine the effect of the "Castle Doctrine" laws, which remove the "duty to retreat" from the home before using lethal force. These are the laws, time period, and models examined by Lott.[29] As in Lott (2010), the resulting estimates imply that these laws decrease violent crime. In particular, the murder rate is estimated to fall by 4%, robbery by 2%, and assault by 10%. Finally, under this specification, SYG laws are estimated to decrease the overall violent crime rate by 6%.

3.2.4 Assessing the invariance assumptions using pre-treatment data

In light of the sensitivity of the estimates to the different invariance models, a natural question to consider is whether a data-driven approach can determine if any of the identifying assumptions are valid. Although the validity of the different invariance models cannot be empirically tested, it has become common to use pre-treatment data to assess the credibility of the assumptions. The basic idea is to compare the pre-treatment estimates of the crime rate implied by the invariance model with the observed crime rate in the thirteen states that adopt SYG laws in 2006. The invariance model, if correct, point-identifies the crime rate that would be realized if the treated states did not have a SYG statute. Pre-2006, the treated states did not have a SYG statute. Thus, if the invariance assumptions are valid in the pre-2006 period, the estimated rate of crime found under the invariance assumption will equal the observed crime rate in the thirteen treated states. If so, this finding may provide some heuristic support in favor of the invariance assumption.

[27] See Cheng & Hoekstra, supra note 3.

[28] See Donohue et al., supra note 5.

[29] See Lott in supra note 3. The treated states include Florida (2005), Illinois (2004), Colorado (1995), New Mexico (1978) , North Carolina (1993), Utah (2003), and Washington (1999).

The crime rates displayed in Figures 1A and 1B provide direct evidence that the invariance assumptions do not hold across all years. All three of the invariance model assumptions are inconsistent with crime rate data pre-2006, before any state in our data adopted a SYG statute. For example, consider the annual murder rates in treated and untreated states displayed in Figure 1A. While these states did not have SYG statutes prior to 2006, the murder rates vary over time and across the groups of states from 1990 to 2006. For example, in 2000, the average murder rate in the thirteen treated states is 1.2 points higher than the average rate in untreated states, nearly four points less than the average murder rate in treated states in 1990, and 0.03 points less than rate found using the DnD invariance model. Thus, the strict invariance assumptions are violated in 2000, before these states adopted SYG statutes. Moreover, the signs and magnitudes of these violations differ over time and across models.

Table 3.7 provides additional evidence on these violations. For each pair of adjacent years from 1990 to 2006, we compute the differences between the observed crime rate in the treated states and the analogous rate found under three different invariance models. In Table 3.7, we display the maximum of the absolute values of these differences. So, for example, the before-after assumption is inconsistent with the observed crime rate in treated states by as much as 85 for violent crime, 1 for murder, 15 for robbery and 60 for assault. Thus, the invariance models evaluated in Section 3.2 are inconsistent with the observed rates in the pre-treatment periods.

While all three of these invariance restrictions are rejected in the pre-2006 years, the existing literature consistently applies these models, especially the DnD and before-after invariance models. Thus, researchers are assuming that the invariance restrictions apply when the outcomes are counterfactual even though they are rejected in periods when the outcomes are observed. While assumptions that are rejected in 2000 may be valid in 2010, the literature using these invariance models does not explain why this might be the case when evaluating SYG laws.

3.3 Bounded variance assumptions

In this section, we use bounded variation assumptions developed in Manski and Pepper to relax the three invariance models evaluated above.[30] For example, rather than assuming the counterfactual murder rate equals the contemporaneous rate from the untreated states, we will instead assume that the counterfactual rate is bounded between the contemporaneous rate plus and minus a constant. A bounded before-after variation assumption restricts the absolute difference in mean treatment response between two periods to be less than some constant. We refer to this constant as the degree of similarity and label it as S. The larger the selected value of the constant, S, the weaker the assumption. These bounded variation assumptions have identifying power because they imply that counterfactual state-year murder rates are similar to observed rates in other states and years.

In Section 3.3.1 we examine the sensitivity of inferences to different identifying restrictions without taking a stand on any particular bounded variation model or the value of the associated degree of similarity, S. Then, in Section 3.3.2, we apply a specific set of bounded variation models that are consistent with the pre- 2006 crime rates. In particular, we use the results in Table 3.7 to define the degree of similarity. Under these bounded variation models, the sign of the average treatment effect is identified for some outcomes, but not others.

3.3.1 Sensitivity of inferences to the bounded variation assumption

We begin by examining the sensitivity of inferences to different bounded variation assumptions without taking a stand on the particular value of the degree of similarity, S. This allows us to illustrate the sensitivity of inferences to different assumptions.

Figure 2 traces out the effect of SYG statutes on average murder rates in the thirteen treated states for different values of S. The traditional contemporaneous invariance assumption, where S = 0, point identifies the average treatment effect, revealing that SYG statutes increase the average murder rate by 1.65. However, ambiguity about the average treatment effect increases with S, and any value of S larger than two renders it impossible to sign the average treatment effect. For example, when $S = 2$, the counterfactual murder rate is estimated to lie between [2.5, 6.5], or the counterfactual murder rate from Table 2 +/- S.

[30] See Manksi & Pepper, supra note 1.

The observed murder rate for the thirteen treated states is 6.2, and the average treatment effect is estimated to lie between [-0.3, 3.7].

Figure 2 also traces out the average treatment effect of SYG under the bounded before-after invariance assumption. The traditional before-after invariance assumption $(S = 0)$ point identifies the average treatment effect, revealing that the SYG statutes decrease the average murder rate by -0.77. Uncertainty about the average treatment effect increases with S, and for values of S in excess of 1 the sign of the average treatment effect is not identified. For example, when $S = 2$, the counterfactual murder rate is estimated to lie between [4.9, 8.9], the observed crime rate for the thirteen treated states is 6.2, and the average treatment effect is estimated to lie between [-2.7, 1.3]. Finally, relaxing the DnD invariance restriction even by small amounts $(S \geq 0.15)$ makes it impossible to sign the average treatment effect.

Rather than consider a single invariance assumption in isolation, it may be sensible to combine different sets of bounded variation assumptions. For example, one might simultaneously assume that the treated and untreated states are similar and that treated states had similar characteristics in 2000-2004 and 2008-2010. Focusing on this joint model, some restrictions on the degree of similarity are rejected while others lead to point identification. For example, it cannot be that the degree of similarity, S, equals zero for both the contemporaneous and before-after bounded variation assumptions. To see this, examine Figure 2 which shows that the point estimate when $S = 0$ for the contemporaneous model differs from the point estimate when S=0 for the before-after model. However, the models point identify the average treatment effect for a variety of parameter values, and identify the sign of the average treatment effect for others. For example, when $S = 1.2$ for both models, the average treatment effect is identified to equal 0.4, and when $S = 0.1$ for the bounded time variation models and 2.3 for the geographic variation models, the average treatment effect is identified to equal -0.7. The average treatment effect is identified to be negative for any feasible values of the degree of similarity such that $S < 0.8$ for before-after models, and greater than zero for any feasible value of the parameters where $S < 1.7$ for contemporaneous models. For other values of the degree of similarity, the bounds do not identify the sign of the average treatment effect.

Overall, this sensitivity analysis traces out the ambiguity resulting from the selection problem. We find that the sign of the average treatment effect is identified for some parameter values, but not others. When different bounded variation models are simultaneously applied, the strict invariance models are ruled out, but other models either point identify the average treatment effect or identify the sign of the average treatment effect. Finally, for some parameter values, the joint model does not identify whether SYG laws increase or decrease murder. Clearly, as with the invariance models, the results are sensitive to the underlying assumptions. In the next section, we evaluate the impact of SYG laws under a set of particular bounded variation models that are based on the pre-2006 data.

3.3.2 Estimates of the effect of SYG laws on bounded variation assumptions

Following Manski and Pepper[31] , we use the pre-treatment period to generate data based degree of similarity parameters, S, for each bounded variation model. In particular, we use the results displayed in Table 3.8 to ensure the bounded variation models are consistent with the observed pre-2006 data. So, for murder, the before-after variation parameter needs to be at least 1.0, the contemporaneous variation parameter at least 1.5 and the DnD variation parameter at least 0.5. For robbery, the analogous parameters are 15, 80, and 15.

Table 3.8 displays results, first considering the three assumptions separately and then combining the assumptions. When the bounded variation models are applied separately, the estimates do not generally identify the sign of the average treatment effect. There are, however, several notable exceptions. For violent crime rates, the average treatment effect is estimated to be positive under the contemporaneous and DnD bounded variation models. Under these models, we estimate that the SYG statutes increase average violent crime rates by at least 2 and by as much as 103. For murder, the average treatment effect is estimated to increase by at least 0.2 and at most 3.2 under the contemporaneous model.

To narrow the bounds, we combine the three bounded variation assumptions. Under this joint bounded variation model, the sign of the average treatment effect is not identified for robbery and assault. For example, the estimated bounds imply that SYG laws might lead the robbery rate to fall by as much as -6 or increase by as much as 11. For the assault rate, the data and models imply that effect of SYG laws lies between -4 and 33. Without stronger assumptions, we cannot determine whether SYG laws increase or decrease the expected rates of robbery or assault.

[31] See Manski & Pepper, supra note 1.

However, under this joint bounded variation model, we estimate that SYG laws increase violent crime and murder. In particular, SYG are estimated to increase the expected murder rate by 0.2, and the violent crime rate by between [3, 50]. This implies that SYG laws increase the murder rate by 3% (from 6.0 to 6.2) and the violent crime rate by at least 1% and as much as 13%. Thus, under this weak bounded variation model, SYG laws are estimated to have a modest positive effect on the average murder and violent crime rates. Cheng and Hoekstra[32], who use a difference-in-difference invariance assumption, draw similar conclusions. In particular, they find SYG laws do not have a statistically significant effect on burglary, robbery, and aggravated assault but a positive effect of 8% on murder.

3.4 Conclusion

Providing credible estimates of the impact of gun laws on crime has proven to be a difficult undertaking. Despite a large empirical literature, research has failed to reach consensus about the impact of different gun laws on crime. Empirical results vary with the data and are highly sensitive to minor variation in the assumptions.

In this paper, we make transparent how assumptions shape inference, focusing on the impact of SYG laws. After illustrating how the empirical findings are sensitive to commonly used invariance assumptions, we then apply the recent methods developed by Manski and Pepper to assess what can be inferred under relatively weak assumptions restricting variation in treatment response across geography and time. These partial identification models highlight the the inherent tradeoff between the strength and credibility of assumptions and findings. Strong invariance assumptions lead to definitive findings that may lack credibility. Weaker bounded variation assumptions can lead to uncertain but credible findings.

By assessing the effect of SYG laws using the bounded variation assumptions, we illustrate the sensitivity of inferences to underlying assumptions. Under the weakest assumptions, the bounds are wide and cannot reveal whether SYG laws increase or decrease violent crimes. Under our preferred joint bounded invariance model, we find evidence SYG laws have uncertain effects on assault and robbery but lead to a modest increase violent crime and murder.

[32] See Cheng & Hoekstra, supra note 3.

3.5 Figures

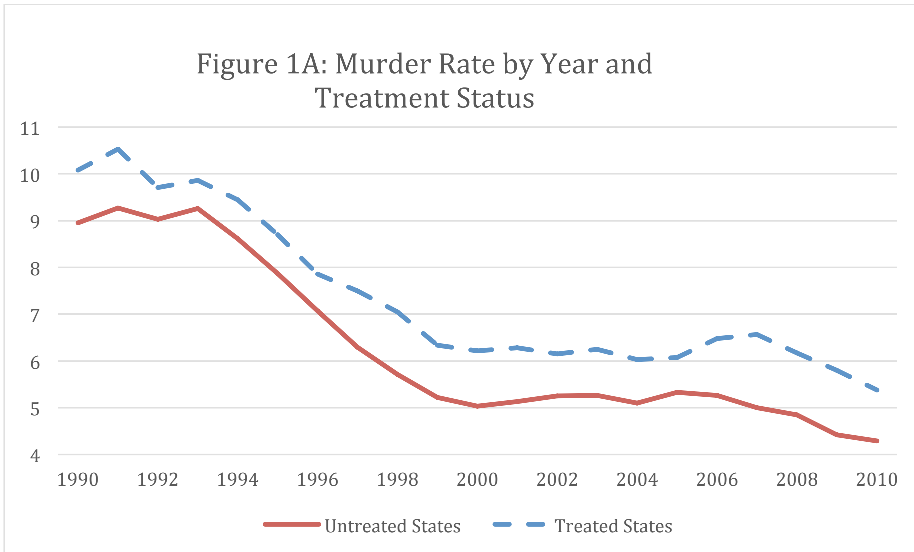
Figure 1A: Murder Rate by Year and
Treatment Status
11
10
9
8
7
6
5
4
1990
1992
1994
1996
1998
2000
2002
2004
2006
2008
2010
Untreated States
Treated States

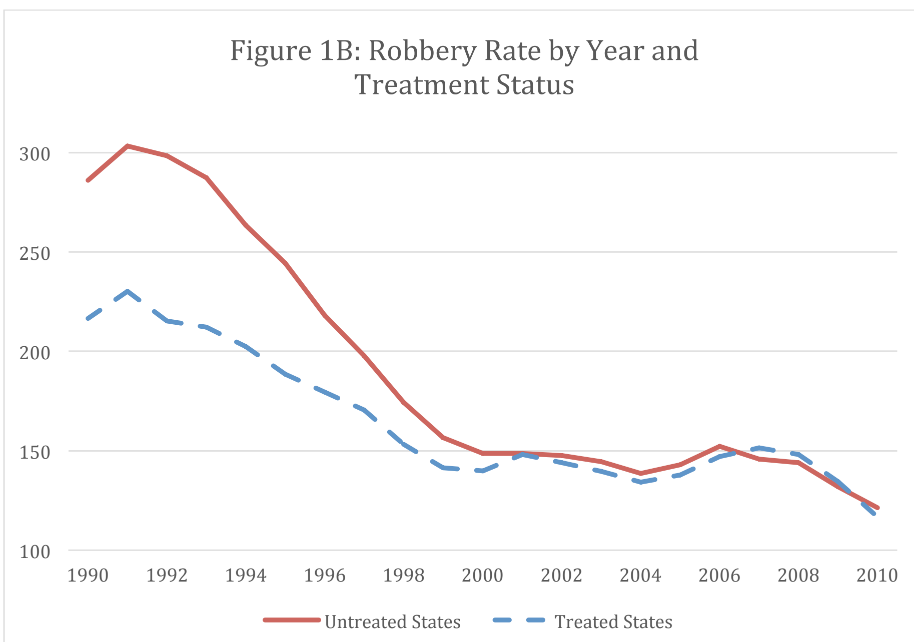
Figure 1B: Robbery Rate by Year and
Treatment Status
300
250
200
150
100
1990
1992
1994
1996
1998
2000
2002
2004
2006
2008
2010
Untreated States
Treated States

Figure 2: Bounds on the ATE of Murder
Under Different Models and S

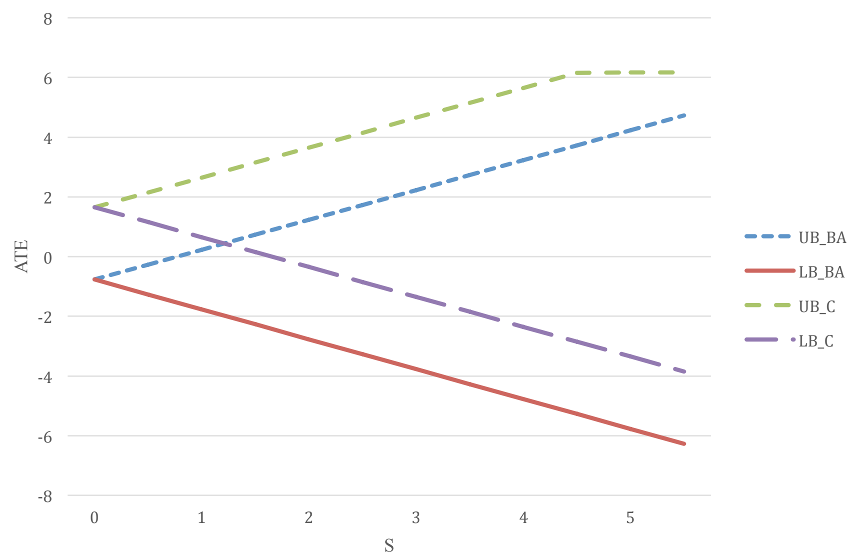

Note: LB ≡ lower bound; UB ≡ upper bound; C ≡ contemporaneous restriction; and BA ≡ before-after restriction.

3.6 Tables

Table 3.1: Murder rates per 100,000 residents by year and treatment status in 2006

	2000-2004	2008-2010
Treated (in 2006)	6.94	6.17
Untreated	5.16	4.52

Note:The treated group includes 13 states. The untreated group includes 29 states and the District of Columbia. The eight states that adopted the Stand Your Ground statuses between 2005 and 2009 (not including 2006) are not included.

Table 3.2: 2008-2010 Murder Rate for the Treated States Under Different Models: Observed, Counterfactual and the Average Treatment Effect (ATE)

	Observed Murder Rate	Counterfactual Murder Rate	ATE
Before-After	6.17	6.94	-0.77
Contemporaneous	6.17	4.52	1.65
DnD	6.17	6.30	-0.12

Table 3.3: Average Treatment Effect Under Different Invariance Models

	Murder	Robbery	Assault	Violent Crime
Before-After	-0.8	-4	-27	-35
Contemporaneous	1.7	-14	56	53
DnD	-0.1	9	18	27

Table 3.4: Linear Panel Model Estimates, 13 Treated States that adopted in 2006

	Murder		Robbery		Assault		Violent crime	
	Log	Level	Log	Level	Log	Level	Log	Level
DnD, Lott (2010) controls	0.09	0.75	0.02	8.31	-0.01	3.68	0.01	12.13
DnD, Donohue (2019) controls	0.12	0.39	0.05	17.15	0.03	-4.11	0.05	13.42
Before/After Lott controls	0.12	0.79	0.09	11.04	-0.01	-3.45	0.03	11.85
Contemp., Lott controls	0.06	0.32	0.21	9.39	0.17	48.13	0.15	60.91

Table 3.5: Linear Panel Data Model Estimates, Cheng and Hoeskstra (2013) Replication, All states that adopted from 2005-2010

	Murder		Robbery		Assault		Violent crime	
	Log	Level	Log	Level	Log	Level	Log	Level
DnD, Lott (2010) controls	0.03	0.13	-0.0001	-2.67	-0.005	-16.06	-0.003	-19.40
DnD, Donohue (2019) controls	0.04	0.10	-0.06	-13.59	-0.02	-34.40	-0.03	-48.01

Table 3.6: Linear Panel Data Model Estimates, Lott (2010) Replication, 1977-2005

	Murder		Robbery		Assault		Violent crime	
	Log	Level	Log	Level	Log	Level	Log	Level
DnD, Lott (2010) controls	-0.04	-0.49	-0.02	5.27	-0.10	-44.37	-0.06	-40.32

Table 3.7: Maximum Absolute Difference between the Observed Annual Crime Rate in the Treated States and the Counterfactual Estimate of the Crime Rate, 1990-2006

	Murder	Robbery	Assault	Violent Crime
Before-After	1.0	15	60	85
Contemporaneous	1.5	80	60	50
DnD	0.5	15	35	25

Table 3.8: Estimated Treatment Effect Under Different Bounded Variation Models

	Murder	Robbery	Assault	Violent Crime
Before-After	[-1.8,0.2]	[-19,11]	[-87,33]	[-120,50]
Contemporaneous	[0.2,3.2]	[-94,67]	[-4,116]	[3,103]
DnD	[-0.6,0.4]	[-6,24]	[-17,53]	[2,52]
DnD+Contemporaneous	[0.2,0.4]	[-6,24]	[-4,53]	[3,52]
DnD+Before-After	[-0.6,0.2]	[-6,11]	[-17,33]	[2,50]
All three models	[0.2,0.2]	[-6,11]	[-4,33]	[3,50]

Chapter 4

Divorce and the Medicaid expansion

4.1 Introduction

In 2014, 26 states enacted the largest federal expansion of Medicaid in decades. Part of the Affordable Care Act, passed in Congress in 2010, the Medicaid expansion provided federal funding to states that extended public health insurance to low-income childless adults and increased access for parents in low-income households. These expansions were associated with an 8 percentage point drop in the share of uninsured individuals nationwide, falling from 20% to 12% (Baumgartner et al 2020) and the increased access to health insurance had wide-ranging effects on beneficiaries.

One potential effect of the Medicaid expansion is on marriage and divorce behavior. Prior to the Affordable Care Act, there were few options for health insurance outside of employer sponsored insurance. Employer sponsored health insurance is the most common source of health insurance coverage in the United States, and these plans often include the option to add a spouse and dependents to the policy. For people receiving coverage through their spouse's employer sponsored plan, divorce leads to a loss of health insurance coverage. As a result, we might expect the expansion of insurance options to increase divorce rates among the set of people receiving coverage through their spouse, because they could now obtain coverage independently of their spouse's plan.

The effects of the Medicaid expansion on marriage and divorce behavior are unlikely to be evenly distributed across states that passed expansions. This is because, prior to the expansion, states differed in the generosity of the program. Some states heavily restricted benefits to adults with children, while others were less strict. Previous work that has studied the effects of the Affordable Care Act on marriage and divorce behavior has used a difference-in-differences estimator (Hampton and Lenhart (2019), Slusky and Ginther (2016)), which conceals all of this variation. Given the limited work studying the relationship between the Medicaid expansions and divorce, it is worth studying whether states that had stricter programs prior to the expansion experienced greater changes in divorce rates after expanding access to Medicaid.

There are three goals to this work. First, I use an alternative dataset to replicate the results from Hampton and Lenhart (2019), who use a difference-in-differences estimation strategy to study the effects of the Medicaid expansion on divorce rates. I do not find statistically significant evidence that the effects seen in Hampton and Lenhart (2019) hold in my dataset. Second, I study whether states that had stricter Medicaid programs prior to the expansion experienced different effects on divorce rates than those that were less strict. I find statistically significant evidence that individuals that live in states that had larger Medicaid expansions were more likely to divorce. Finally, I also check for heterogeneous effects of the expansion and find that individuals between ages 27 and 49 and individuals with a high school degree or below experienced the strongest effects from the expansions.

4.2 Background

4.2.1 Medicaid expansions

Medicaid provides free health insurance to low income households. In order to qualify for benefits, households must meet income eligibility thresholds, and if they earn above the threshold, they are ineligible for benefits. These thresholds are generally expressed as a percentage of the federal poverty line (FPL), which varies by household size. The thresholds also vary by the category of beneficiary. Children, disabled adults, and parents in low-income households all face different eligibility thresholds.

A key component of the Medicaid expansion was a large change in income eligibility thresholds. Prior to the expansions in 2014, states independently set income eligibility thresholds, although childless adults were ineligible for benefits in all states. The Medicaid expansions of 2014 set the thresholds for childless adults and parents in low-income households to 138% of FPL, for all states that adopted the expansions.

States that adopted the Medicaid expansions varied in their income eligibility thresholds for parents in 2013, prior to the expansions. This means that, among states that adopted expansions, the size of the change in eligibility thresholds for parents varied. As an example, in Arkansas, which expanded Medicaid, the income eligibility thresholds for parents to receive benefits was 16% of FPL in 2013, and increased to 138% in 2014. For households of size 4, this meant that in order for parents to receive Medicaid in 2013, annual household earnings could not exceed $3,768. In 2014, this limit increased to $45,420. This is a very large increase in the number of parents that were newly eligible for coverage. On the other hand, Ohio increased the income eligibility threshold for parents from 96% of FPL to 138% of FPL, which constitutes a much smaller change in the number of newly eligible parents. This variation in the number of newly eligible parents implies that any effects observed from Medicaid expansions are likely to vary among the states that passed expansions.

4.2.2 Relationship between marriage and health insurance

One potential effect of the Medicaid expansions is on marriage and divorce behavior. The theory of the relationship between marriage and health insurance is primarily based on the model of marriage suggested by Becker (1981). In this work, divorce occurs when the utility from remaining in the marriage is less than the utility from leaving the marriage. Losing access to health insurance through the spouse's coverage is hypothesized to reduce utility upon divorce, and so employer sponsored health insurance can create "marriage lock" to avoid the loss of health insurance. Under this theory, the Medicaid expansion would be expected to increase the divorce rate, as divorcees would have access to health insurance through Medicaid.

The effect of health insurance on marriage decisions has been studied previously in the literature. For example, Yelowitz (1998) shows that the extension of Medicaid to married parents increased the probability of marriage by around 1.7 percentage points. Chen (2019) shows that the Medicare eligibility threshold at age 65 is associated with a 7 percentage point increase in the probability of divorce once individuals become eligible. Both of these papers show that the introduction of public health insurance coverage affects marriage and divorce behavior.

There has also been previous work looking at the effects of the Medicaid expansions on marriage and divorce behavior. Hampton and Lenhart (2019) study the effect of the Medicaid expansions on divorce using a traditional difference in differences estimation strategy. In this working paper, the authors use a difference-in-differences identification strategy with Current Population Survey (CPS) data. They find that the Medicaid expansions are associated with a 2% decrease in the probability of marriage and a 4% increase in the probability of divorce. This is the work that I build upon, both by replicating their work and using alternate sources of variation.

Slusky and Ginther (2017) find that the Medicaid expansions decreased the likelihood of divorce among older couples. The hypothesis in their paper is that “medical” divorces occurred prior to the Medicaid expansions to provide coverage to one spouse with high healthcare costs. By increasing the income eligibility thresholds and removing asset tests, the Medicaid expansions eliminated the need for some couples to divorce to gain coverage.

Abramowitz (2016) studies the effect of the young adult coverage mandate portion of the Affordable Care Act on the marital behavior of young adults. The Affordable Care Act extended coverage to young adults under their parent’s health insurance plan until age 26. The author finds that this extension of coverage is associated with a decline in marriage, a decline in cohabitation, and an increase in the probability of divorce.

4.3 Data

To study the effects of the Medicaid expansions on divorce, I use American Community Survey (ACS) data from 2011-2017. This provides 3 years of data prior to the expansions and 4 years of post expansion data. The ACS is an annual cross-sectional survey of around 3.5 million addresses. I construct the dataset following Hampton and Lenhart (2019).

My interest is in studying marriage and divorce behavior, so I restrict my dataset to individuals between 18 and 64 years old. Individuals under 18 in the ACS are very unlikely to be married and even less likely to be divorced, so I exclude them from the analysis. Adults over 64 are eligible for Medicare, and thus unlikely to respond to Medicaid eligibility changes.

I also exclude from my final dataset individuals living in the 17 states that had enacted a Medicaid expansion prior to 2014. Prior to the ACA, states could apply for waivers to make changes to their Medicaid programs. These waivers could include a variety of different policy instruments, like co-pays for health care, increased income eligibility thresholds, and premiums. Because these waiver programs vary widely and likely affect how individuals in the state respond to the expansion, I exclude them from my dataset.

In Table (4.1), I include summary statistics for states that enacted stand-alone Medicaid expansions in 2014 and those that did not. I end up with around 7.5 million individuals over 7 years, after excluding everyone under 18 and above 64 and individuals in states that had enacted a Medicaid expansion waiver prior to 2014. Looking specifically at divorce, in my sample, across all 7 years, the average share of individuals who were divorced was 0.12 in states that did not pass expansions and 0.13 in those that did, and the share of individuals who were divorced in the last year was 0.01 in both groups. In Figures (4.1) and (4.2), I also include graphs showing trends in divorce rates, comparing states that enacted Medicaid expansions in 2014 and those that didn't. The graphs show that states that enacted stand-alone expansions in 2014 had consistently higher rates of divorced individuals, while the rate of individuals who were divorced in the last year does not appear to differ between the two groups.

4.4 Empirical strategy

The first strategy that I use to estimate the effect of the Medicaid expansions on divorce rates is difference in differences. In the treatment group, I include the 9 states that introduced a Medicaid expansion in 2014 and had no previous expansion, and in the control group, I include the 24 states that did not pass an expansion in 2014.[1].

The first model that I estimate is:

$$100 * y_{ist} = \alpha_0 + \alpha_1 * Treat_s + \alpha_2 * Post_t * Treat_s + \alpha_3 * X_{ist} + \gamma_t + \psi_s + \epsilon_{ist} \quad (4.1)$$

In this equation, $Post$ is one if t is greater than 2013, while $Treat$ is equal to one if individual i lives in one of the 9 states that introduced a stand-alone Medicaid expansion in 2014. I also include a vector of control variables in X, including employment status, education, gender, age, and race, and year and state fixed effects. I include year fixed effects to account for the general negative trend in divorce rates, and state fixed effects to account for underlying differences in divorce rates across states. Control variables included in X

[1]The treatment states are: Arkansas, Kentucky, Michigan, Nevada, New Hampshire, New Mexico, North Dakota, Ohio, and West Virginia

are meant to account for the fact that some subpopulations are more likely to be eligible for Medicaid, and so more likely to experience the effects of the Medicaid expansion. Standard errors are clustered at the state level.

I consider two outcomes, y_{ist}. The first outcome is whether the individual is divorced, and y_{ist} equals 1 if individual i living in state s at time t is divorced and 0 otherwise. The second outcome I consider is whether the individual divorced in the last year, and y_{ist} equals 1 if individual i living in state s at time t reports that they got divorced within the last 12 months. This is a better measure of divorces that are "caused" by the Medicaid expansions, because it excludes divorces that happened decades ago and are unrelated to the expansions.

In equation (4.1), the coefficient of interest is α_2. I multiply my outcome variable by 100 so that the estimated coefficients are easier to read. This means that the interpretation of α_2 is the percentage point change in the probability of divorce post-2013 due to the Medicaid expansion.

The effects of the Medicaid expansion on divorce may increase over time because divorces can take years to process through the courts. To study whether this is true in my dataset, I estimate the following model:

$$100 * y_{ist} = \alpha_0 + \alpha_1 * Treat_s + \Sigma_{j=2012}^{j=2017} \beta_j * Year_j * Treat_s + \alpha_2 * X_{ist} + \gamma_t + \psi_s + \epsilon_{ist} \quad (4.2)$$

This regression allows the effect of the Medicaid expansion to vary over time and serves two purposes. First, for years after 2013, we can study whether the effect of the Medicaid expansion on divorce increases over time. Second, for years before 2014, we can look for evidence that the parallel trends assumption holds. The parallel trends assumption says that states that passed the expansions would have experienced the same change in divorce rates after 2013 as states that did not pass the expansion, if they had not passed the expansion. One commonly applied test of this assumption is to test that states that passed the expansion had the same trends in divorce rate as states that did not pass the expansions pre-2014. In equation (4.2), this would imply that β_{2012} and β_{2013} are statistically insignificant. The two regressions above serve to verify that trends observed in Hampton and Lenhart (2019) hold in the ACS, which has a larger sample than the CPS.

However, the effects of the Medicaid expansion are unlikely to be evenly distributed across the states that introduced the expansions. In particular, some states increased the income thresholds for parents more than other states, and a larger increase is likely to affect more couples. As a result, I construct a "treatment exposure" variable, that is the difference between the January 2013 income eligibility threshold for parents, and the threshold in January 2014, which was 1.38, or 138% of the FPL. Table (4.2) lists the treatment exposure for each treatment state and the average treatment exposure among treatment states, which was 0.78.

The regression that I run is:

$$100 * y_{ist} = \alpha_0 + \alpha_1 * Treat_s + \alpha_2 * Post_t * Treat_s \\ + \alpha_3 * Post_t * Treat_s * Exp_s + \alpha_4 * X_{ist} + \gamma_t + \psi_s + \epsilon_{ist} \quad (4.3)$$

In this regression, the interesting coefficients are α_2 and α_3, which measure both how the probability of divorce varies in states that passed expansions and how the probability of divorce varies with the size of the change in income eligibility thresholds for parents. The estimated effect of the Medicaid expansion on the divorce rate in state s with treatment exposure Exp_s is equal to $\alpha_2 + \alpha_3 * Exp_s$.

Finally, I also study whether the effect of exposure to the Medicaid expansions on divorce rates varies over time. To do this, I estimate the following model:

$$100 * y_{ist} = \alpha_0 + \alpha_1 * Treat_s + \Sigma_{j=2012}^{j=2017} \beta_j * Year_j * Treat_s + \alpha_2 * Treat_s * Exp_s \\ + \Sigma_{j=2012}^{j=2017} \tau_j * Year_j * Treat_s * Exp_s + \alpha_3 * X_{ist} + \gamma_t + \psi_s + \epsilon_{ist} \quad (4.4)$$

As in equation (4.2), equation (4.4) measures both how the effects of exposure to the Medicaid expansions varied post expansion, and verifies that the parallel trends assumption holds within the treatment group. If the parallel trends assumption holds within the treatment group, we would expect to see that τ_{2012} and τ_{2013} are statistically insignificant, suggesting that states with lower income eligibility thresholds prior to the expansion do not have different underlying trends in divorce rates than those with higher income eligibility thresholds prior to the expansion.

4.5 Results

4.5.1 Main results

First, I replicate the results of Hampton and Lenhart (2019) using equation (4.1). These results are reported in columns (1) and (2) of Table (4.3). The treatment effect is the coefficient of $Post * Treat$. In column (2), I estimate a treatment effect of 0.077, which means that Medicaid expansions led to an increase in the percent of divorced adults of 0.077 percentage points. The average divorce rate in the treatment states before the expansions was 0.13 or 13%, and so this constitutes an increase of $\frac{0.077}{13} = 0.05\%$, because the dependent variable is multiplied by 100. This result is statistically insignificant. These results differ from the results reported in Hampton and Lenhart (2019), who find that the Medicaid expansions increased divorce rates by around 4% using this estimation strategy with CPS data. I also find that divorce rates increase in states that passed a Medicaid expansion, but the magnitude and statistical significance of my estimates differ, because I am using ACS data rather than CPS data.

In the third column of Table (4.3), I present results from estimating equation (4.3), introducing treatment exposure to the traditional difference-in-differences model. From this regression, it is clear that the effect of the Medicaid expansion on divorce rates varies with the size of the change in the parents income eligibility threshold. To calculate the total treatment effect estimated in this regression, the formula is $-0.428 + 0.74 * Exp$. The average treatment exposure among treatment states was 0.78, and the predicted effect of the Medicaid expansion on divorce rates is positive for most treatment states. As an example, these estimates imply that the predicted increase in the percentage of divorced individuals living in Michigan, which had a treatment exposure of 0.74, is $-0.428 + 0.74 * 0.74 = 0.12$ percentage points, conditional on covariates. This is a 1% increase in the share of divorced individuals over the pre-treatment average, which was around 0.125, or 12.5%.

In Table (4.4), I present results using whether the individual has been divorced in the last year as the dependent variable. Column (1) presents the traditional difference-in-differences estimates from equation (4.1), while column (2) presents estimates introducing the treatment exposure variable from equation (4.3). Using divorced in the last year as the dependent variable, I again find statistically significant evidence that states that experienced large increases in the income eligibility threshold for parents experienced larger increases in divorce rates than those that had smaller increases. These results are reported in column (2). To interpret this coefficient, again, take the example of Michigan. These results suggest that the Medicaid expansion in Michigan increased the percent of individuals that divorced within the last year by $-0.054 + 0.114 * 0.74 = 0.03$ percentage points. This is a 3% increase over the percent of individuals that divorced within the last year in 2013 among treatment states, which was 1.1%.

4.5.2 Effects over time

I present estimates of equations (4.2) and (4.4) in Table (4.5), which both verifies that the parallel trends assumption holds, and estimates the effect of the expansions over time. I present results using both whether the individual is divorced and whether the individual divorced within the last year as dependent variables.

First, discussing the parallel trends assumption, columns (1) and (3) show that for both dependent variables, prior to the Medicaid expansions, the treatment group satisfies the parallel trend assumption, as neither $Treat*2012$ nor $Treat*2013$ are statistically significant. The results from estimating equation (4.4), in column (4), show that, for the dependent variable divorced in the last year, trends in divorce rates among treatment states did not vary by Exp prior to the expansion. In this column, the interaction of treatment exposure and year is statistically insignificant prior to 2017. This suggests that, within the treatment group, states with larger increases in income eligibility thresholds in 2014 did not have different trends in divorce rates than those with smaller increases prior to the expansions. In column (2), there is a statistically significant relationship between treatment exposure and divorce rates in 2013, suggesting that the parallel trends assumption may not be satisfied within the treatment group using divorced as the dependent variable.

Next, discussing the effects of the Medicaid expansions over time, column (1) of Table (4.5) suggests that there is no effect on the divorce rate in any year from the expansions, while column (2) reveals that states that were more exposed to the expansion experienced increased divorce rates in 2016 and 2017. In column (4), it also appears that the effects of the expansion on the share of individuals divorced in the last year increased over time. By 2017, the Medicaid expansions led to higher divorce rates in states with more exposure. Table (4.5) shows that the positive and statistically significant relationship between the Medicaid expansions and divorce rates in Tables (4.3) and (4.4) is primarily driven by divorce rates in 2017, 3 years after the expansions.

4.5.3 Heterogeneous effects

I check for heterogeneous effects from the expansions by education, age and whether there are children in the household. To perform this analysis, I run equations (4.1) and (4.3) by subgroups. In these regressions, I use whether the individual is divorced in the last year as the outcome variable.

We would expect less educated individuals to be more affected by the Medicaid expansions because they are more likely to be income-eligible. We would also expect households with children to be more responsive to treatment exposure, because treatment exposure measures the change in the income eligibility threshold for parents. Finally, we would expect to see the strongest effects among individuals in the 27 to 39 and 40 to 49 age ranges. Individuals between 18 and 26 can still be covered under their parent's policy, and so already had an outside option for health insurance that is not related to marriage. As mentioned in Section 2, older individuals experienced a decline in "medical divorces" following the expansions (Slusky and Ginther (2016)), and so we would expect the divorce rate to decline among individuals in this age range.

Starting with the Children in HH panel in Table (4.6), there is evidence in column (1) that individuals with children are more responsive to the Medicaid expansions, with individuals with no children being less likely to divorce after an expansion. In column (2), the predicted effect of the Medicaid expansion on the share of individuals divorced last year is again negative or very near zero for most states for adults with no children in the household. For instance, in Michigan, the probability that an individual with no children in the household was divorced in the last year is predicted to increase by $-0.86 + 0.123 * 0.74 = 0.005$ percentage points due to the Medicaid expansion. Greater exposure to the expansion is also predicted to increase the probability that an individual with children is divorced in the last year, but this result is statistically insignificant.

In the second panel of Table (4.6), individuals in the middle two age categories are more affected by the Medicaid expansions. Column (2) shows that individuals in all age categories are more responsive to the Medicaid expansions in states where the changes in income eligibility thresholds were larger. Taking again the example of Michigan, the Medicaid expansion is estimated to increase the probability that an individual between 18 and 26 years old is divorced in the last year by $-0.07 + 0.138 * 0.74 = 0.03$. The average divorce rate pre-expansion in the treatment group in this age category was 0.005, or 0.5%, so an increase of 0.03 percentage points is a 6% increase in the share of individuals divorced last year. Among individuals between 27 and 39, the share of individuals divorced in the last year was predicted to increase by 3% post-expansion, and by 2% among individuals between 40 and 49. While the effect is largest among younger individuals, the effect is only marginally significant for the youngest age group.

The last panel of Table (4.6) shows how the treatment effect varies by educational attainment. These results indicate that less educated individuals are more responsive to the Medicaid expansions, and are more likely to divorce in states with larger increases in the Medicaid income eligibility threshold. The largest estimated effect of the expansion is among individuals with a high school diploma or less. Individuals with a high school diploma or less were both more likely to be divorced in the last year across all expansion states and more responsive to the size of the change of the income eligibility thresholds.

4.6 Conclusion

This paper shows how the effects of the Medicaid expansion on divorce rates varies by exposure to the expansions. When comparing all states that passed a stand-alone Medicaid expansion in 2014 with states that did not pass an expansion, I find no evidence that Medicaid expansions increased divorce rates. After introducing a treatment exposure variable, which measures the change in income eligibility thresholds for parents coverage between 2013 and 2014, I find that states that enacted larger expansions experienced increased divorce rates of around 1%. I also find that the share of individuals that reported a divorce within the last year increased by around 3% following expansions and that this increase is larger in states that passed larger expansions.

4.7 Tables

Table 4.1: Summary statistics

	No expansion	Medicaid expansion in 2014
Divorced	0.12	0.13
Divorced last year	0.01	0.01
Share white	0.77	0.84
Share black	0.14	0.09
Average age	41.96	42.42
Share in labor force	0.73	0.73
Share with high school diploma or more	0.65	0.63
Share with children in HH	0.39	0.37
N	6,118,903	1,588,256

No expansion: 24 that states did not pass an expansion in 2014

Medicaid expansion in 2014: 9 states that passed a stand-alone expansion in 2014

Excludes 17 states that passed an expansion prior to 2014 and individuals less than 18 and over 64

Table 4.2: Treatment exposure

	Income threshold Jan. 2013	Income threshold Jan. 2014	Exposure
Arkansas	0.16	1.38	1.22
Kentucky	0.57	1.38	0.81
Michigan	0.64	1.38	0.74
Nevada	0.84	1.38	0.54
New Hampshire	0.47	1.38	0.91
New Mexico	0.85	1.38	0.53
North Dakota	0.57	1.38	0.81
Ohio	0.96	1.38	0.42
West Virginia	0.31	1.38	1.07
Average exposure			0.78

Exposure: Income threshold in 2014-Income threshold in 2013

Table 4.3: Replication

VARIABLES	(1) Divorced	(2) Divorced	(3) Divorced	Mean divorce rate pre-2014
Treat	0.894***	0.852***	0.691***	12.85
	(0.049)	(0.080)	(0.096)	
Post*Treat	0.039	0.077	-0.428**	
	(0.087)	(0.090)	(0.158)	
Post*Treat*Exposure			0.740***	
			(0.253)	
Additional controls	No	Yes	Yes	
Year FE	Yes	Yes	Yes	
State FE	Yes	Yes	Yes	
Observations	7,707,159	7,707,159	7,707,159	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Additional controls: Age, presence of children in hh, gender, employment status, education and race

Standard errors clustered at state level

Estimates of equations (4.1) and (4.3), using divorced as dependent variable

Table 4.4: Divorced last year

VARIABLES	(1) Divorced last year	(2) Divorced last year	Mean divorce rate pre-2014
Treat	0.031**	0.006	1.16
	(0.012)	(0.013)	
Post*Treat	0.024	-0.054	
	(0.020)	(0.032)	
Post*Treat*Exposure		0.114***	
		(0.036)	
Additional controls	Yes	Yes	
Year FE	Yes	Yes	
State FE	Yes	Yes	
Observations	7,707,159	7,707,159	

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Additional controls:Age, presence of children in hh, gender, employment status, education and race

Standard errors clustered at state level

Estimates of equations (4.1) and (4.3), using divorced last year as dependent variable

Table 4.5: Effect over time

VARIABLES	(1) Divorced	(2) Divorced	(3) Divorced last yr	(4) Divorced last yr
Treat*2012	0.075	0.383*	0.064	0.075
	(0.103)	(0.214)	(0.040)	(0.089)
Treat*2013	-0.122	-0.424**	0.029	0.020
	(0.123)	(0.161)	(0.042)	(0.071)
Treat*2014	-0.066	-0.259	0.060*	-0.015
	(0.114)	(0.231)	(0.031)	(0.095)
Treat*2015	0.031	-0.174	0.075*	0.059
	(0.101)	(0.161)	(0.043)	(0.078)
Treat*2016	0.149	-0.652***	0.018	-0.069
	(0.170)	(0.194)	(0.046)	(0.092)
Treat*2017	0.132	-0.689***	0.066*	-0.064
	(0.166)	(0.153)	(0.033)	(0.068)
Treat*Exposure*2012		-0.450		-0.017
		(0.370)		(0.135)
Treat*Exposure*2013		0.443*		0.013
		(0.227)		(0.109)
Treat*Exposure*2014		0.282		0.111
		(0.298)		(0.111)
Treat*Exposure*2015		0.300		0.023
		(0.219)		(0.113)
Treat*Exposure*2016		1.174***		0.127
		(0.288)		(0.122)
Treat*Exposure*2017		1.202***		0.190**
		(0.151)		(0.086)
Additional controls	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Observations	7,707,159	7,707,159	7,707,159	7,707,159

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Additional controls:Age, presence of children in hh, gender, employment status, education and race

Standard errors clustered at state level

Estimates of equations (4.2) and (4.4), showing treatment effects over time

Table 4.6: Heterogeneous effects

		Divorced last yr (1)	Divorced last yr (2)	Mean pre-2014
No children in HH	Post*Treat	-0.002	-0.086***	1.10
		(0.018)	(0.021)	
	Post*Treat*Exp		0.123***	
			(0.027)	
Children in HH	Post*Treat	0.068*	0.001	1.24
		(0.035)	(0.069)	
	Post*Treat*Exp		0.098	
			(0.073)	
Age 18 to 26	Post*Treat	0.024	-0.070	0.50
		(0.032)	(0.047)	
	Post*Treat*Exp		0.138*	
			(0.073)	
Age 27 to 39	Post*Treat	0.026	-0.191**	1.67
		(0.038)	(0.078)	
	Post*Treat*Exp		0.320***	
			(0.106)	
Age 40 to 49	Post*Treat	0.023	-0.117	1.61
		(0.034)	(0.081)	
	Post*Treat*Exp		0.205**	
			(0.089)	
Age 50 to 64	Post*Treat	0.023	0.076**	0.88
		(0.031)	(0.031)	
	Post*Treat*Exp		-0.078**	
			(0.036)	
HS diploma	Post*Treat	0.060*	-0.071	1.12
		(0.032)	(0.070)	
	Post*Treat*Exp		0.190**	
			(0.071)	
Assoc. deg.	Post*Treat	0.007	-0.076*	1.26
		(0.022)	(0.039)	
	Post*Treat*Exp		0.122**	
			(0.056)	
Bachelor's deg.	Post*Treat	-0.010	-0.007	1.01
		(0.042)	(0.070)	
	Post*Treat*Exp		-0.005	
			(0.101)	
Post grad	Post*Treat	-0.024	0.101	0.92
		(0.037)	(0.104)	
	Post*Treat*Exp		-0.188	
			(0.125)	

Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Estimates of equations (4.1) and (4.3) by children in the household, age, and educational attainment; regressions include controls, state and year fe, with std errors clustered at state level

4.8 Figures

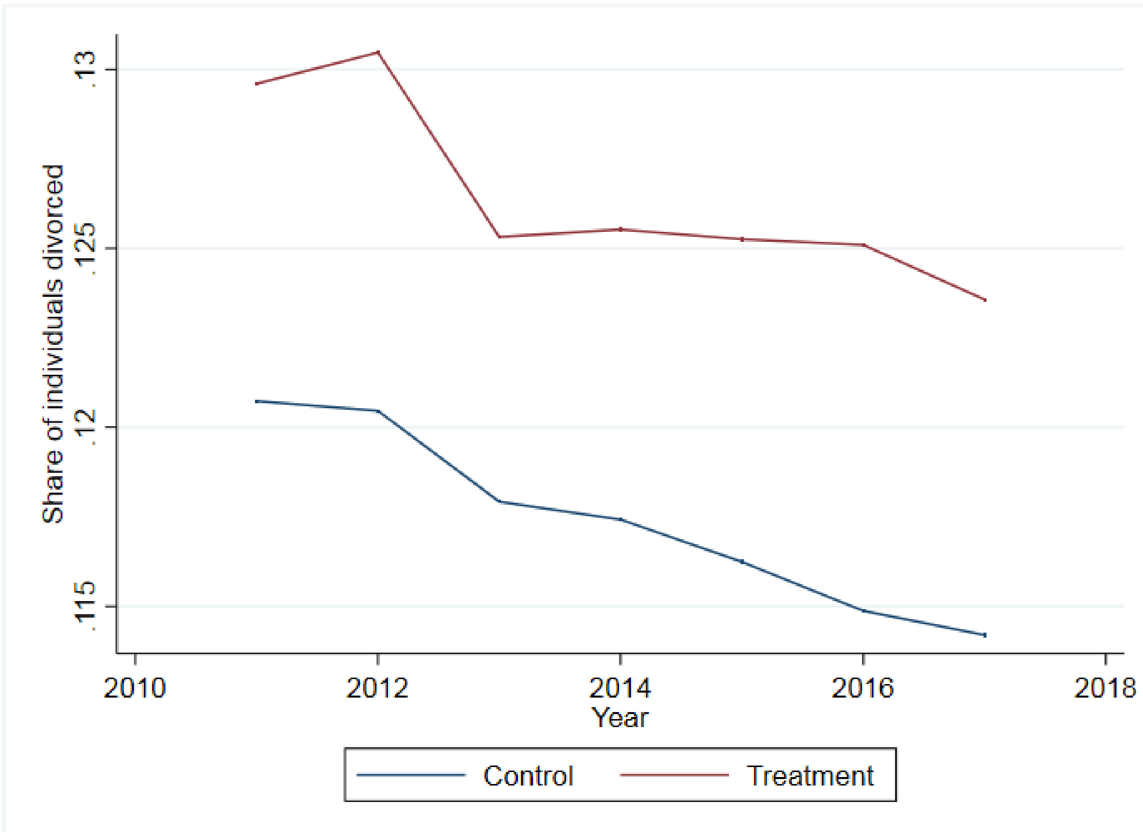

Figure 4.1: Share divorced

Treatment:9 states that passed a stand-alone expansion in 2014; control: 24 states that passed no expansion in 2014

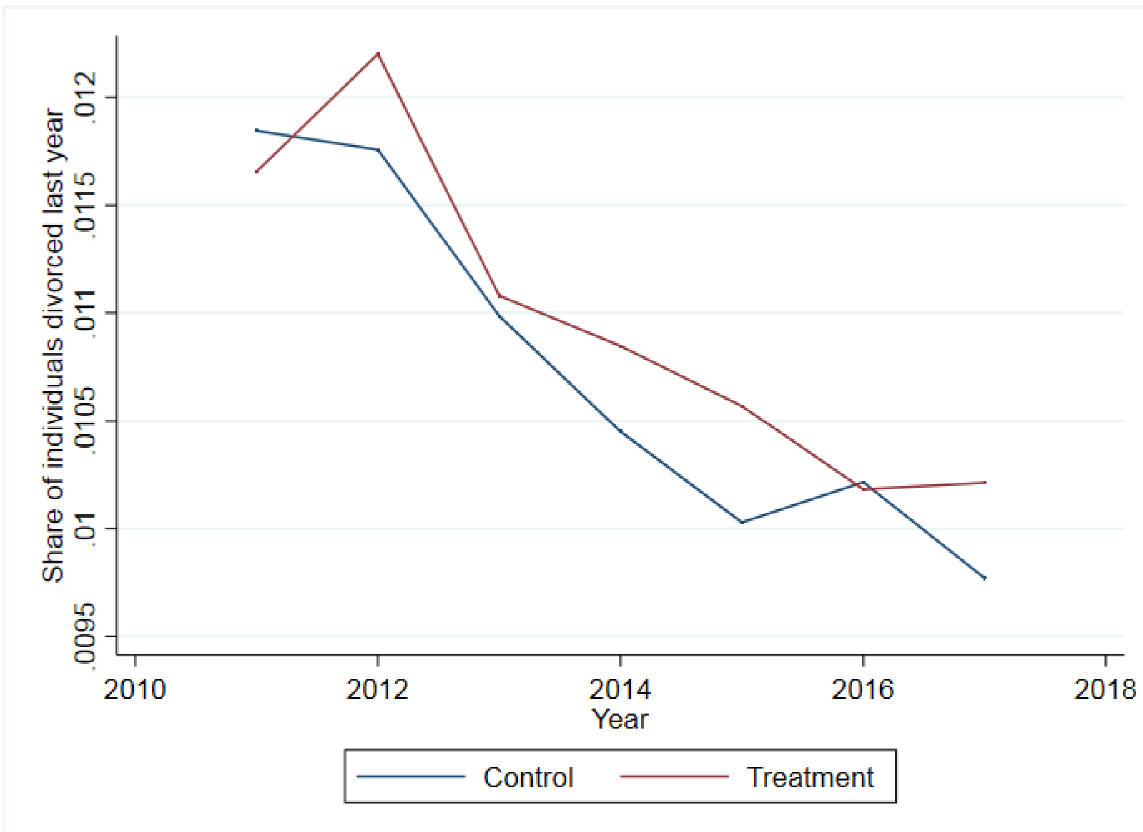

Figure 4.2: Share divorced last year

Treatment:9 states that passed a stand-alone expansion in 2014; control: 24 states that passed no expansion in 2014

Bibliography

[1] Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge. When should you adjust standard errors for clustering? Working Paper 24003, National Bureau of Economic Research, November 2017.

[2] Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. Finite population causal standard errors. Working Paper 20325, National Bureau of Economic Research, July 2014.

[3] Joelle Abramowitz. Saying, "I don't": The effect of the Affordable Care Act young adult provision on marriage. *Journal of Human Resources*, 51(4):933–960, 2016.

[4] Abhay Aneja, III Donohue, John J., and Alexandria Zhang. The Impact of Right-to-Carry Laws and the NRC Report: Lessons for the Empirical Evaluation of Law and Policy. *American Law and Economics Review*, 13(2):565–631, 10 2011.

[5] Jesse C. Baumgartner, Sara R. Collins, David C. Radley, and Susan L. Hayes. How the Affordable Care Act has narrowed racial and ethnic disparities in access to health care. Working papers, Commonwealth Fund, January 2020.

[6] Gary S. Becker. *A Treatise on the Family*. Number beck81-1 in NBER Books. National Bureau of Economic Research, Inc, January 1981.

[7] Adriana Camacho and Emily Conover. Manipulation of social program eligibility. *American Economic Journal: Economic Policy*, 3(2):41–65, 2011.

[8] Mitchell B. Chamlin. An assessment of the intended and unintended consequences of arizona's self-defense, home protection act. *Journal of Crime and Justice*, 37(3):327–338, 2014.

[9] Mitchell B. Chamlin and Andrea E. Krajewski. Use of force and home safety: An impact assessment of oklahoma's stand your ground law. *Deviant Behavior*, 37(3):237–245, 2016.

[10] Tianxu Chen. Health Insurance Coverage and Marriage Behavior: Is There Evidence of Marriage Lock? Working papers 2019-09, University of Connecticut, Department of Economics, May 2019.

[11] Cheng Cheng and Mark Hoekstra. Does strengthening self-defense law deter crime or escalate violence? Evidence from expansions to castle doctrine. *Journal of Human Resources*, 48(3):821–854, 2013.

[12] Raj Chetty, John N. Friedman, Tore Olsen, and Luigi Pistaferri. Adjustment costs, firm responses, and labor supply elasticities: Evidence from Danish tax records. *Quarterly Journal of Economics*, 126:749–804, 2011.

[13] Janet Currie and Jonathan Gruber. Health Insurance Eligibility, Utilization of Medical Care, and Child Health*. *The Quarterly Journal of Economics*, 111(2):431–466, 05 1996.

[14] John J. Donohue, Abhay Aneja, and Kyle D. Weber. Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis. *Journal of Empirical Legal Studies*, 16(2):198–247, 2019.

[15] Vincent Ferraro and Saran Ghatak. Expanding the castle: Explaining stand your ground legislation in american states, 2005–2012. *Sociological Perspectives*, 62(6):907–928, 2019.

[16] Amy Finkelstein, Nathaniel Hendren, and Erzo FP Luttmer. Forthcoming. the value of medicaid: Interpreting results from the oregon health insurance experiment. *Journal of Political Economy*.

[17] Trevor S. Gallen. Using participant behavior to measure the value of social programs: The case of Medicaid. *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, 108:1–54, January 2015.

[18] Craig Garthwaite, Tal Gross, and Matthew J. Notowidigdo. Public health insurance, labor supply, and employment lock. *Quarterly Journal of Economics*, pages 653–696, 2014.

[19] Alexander M. Gelber, Damon Jones, and Daniel W. Sacks. Earnings adjustment frictions: Evidence from the Social Security Earnings Test. NBER Working Paper, August 2014.

[20] Mark Gius. The relationship between stand-your-ground laws and crime: A state-level analysis. *The Social Science Journal*, 53(3):329–338, 2016.

[21] Tal Gross and Matthew J. Notowidigdo. Health insurance and the consumer bankruptcy decision: Evidence from expansions of medicaid. *Journal of Public Economics*, 95(7):767 – 778, 2011.

[22] Matt Hampton and Otto Lenhart. The effect of the ACA Medicaid expansion on marriage behavior. SSRN, September 2019.

[23] Nathaniel Hendren. The policy elasticity. *Tax Policy and the Economy*, 30(1):51–89, 2016.

[24] Hilary W. Hoynes and Diane Whitmore Schanzenbach. Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Applied Economics*, 1(4):109–39, October 2009.

[25] Luojia Hu, Robert Kaestner, Bhashkar Mazumder, Sarah Miller, and Ashley Wong. The effect of the patient protection and affordable care act medicaid expansions on financial wellbeing. Working Paper 22170, National Bureau of Economic Research, April 2016.

[26] David K. Humphreys, Antonio Gasparrini, and Douglas J. Wiebe. Evaluating the Impact of Florida's "Stand Your Ground" Self-defense Law on Homicide and Suicide by Firearm: An Interrupted Time Series Study. *JAMA Internal Medicine*, 177(1):44–50, 01 2017.

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

[28] Henrik J. Kleven and Mazhar Waseem. Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan *. *The Quarterly Journal of Economics*, 128(2):669, 2013.

[29] Henrik Jacobsen Kleven. Bunching. *Annual Review of Economics*, 8:435–464, 2016.

[30] Brent Kreider, John V. Pepper, Craig Gundersen, and Dean Jolliffe. Identifying the effects of snap (food stamps) on child health outcomes when participation is endogenous and misreported. *Journal of the American Statistical Association*, 107(499):958–975, 2012.

[31] John R. Lott. *More Guns, Less Crime: Understanding Crime and Gun-Control Laws*. University of Chicago Press, 2010.

[32] Charles F. Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, 80(2):319–323, 1990.

[33] Charles F. Manski. *Partial Identification of Probability Distributions*. Springer-Verlag, 2003.

[34] Charles F. Manski. *Identification for Prediction and Decision*. Harvard University Press, 2007.

[35] Charles F. Manski and Daniel S. Nagin. Bounding disagreements about treatment effects: A case study of sentencing and recidivism. *Sociological Methodology*, 28(1):99–137, 1998.

[36] Charles F. Manski and John V. Pepper. Monotone instrumental variables: With an application to the returns to schooling. *Econometrica*, 68(4):997–1010, 2000.

[37] Charles F. Manski and John V. Pepper. Deterrence and the death penalty: Partial identification analysis using repeated cross sections. *Journal of Quantitative Criminology*, 29(1):123–141, 2013.

[38] Charles F. Manski and John V. Pepper. How do right-to-carry laws affect crime rates? Coping with ambiguity using bounded-variation assumptions. *Review of Economics and Statistics*, 100(2):232–244, 2018.

[39] Chandler B. McClellan and Erda Tekin. Stand your ground laws, homicides and injuries. *Journal of Human Resources*, 52(3):621–653, 2017.

[40] Bruce D. Meyer and Dan T. Rosenbaum. Welfare, the Earned Income Tax Credit, and the labor supply of single mothers. *The Quarterly Journal of Economics*, 116(3):1063–1114, 2001.

[41] John V. Pepper. The intergenerational transmission of welfare receipt. *The Review of Economics and Statistics*, 82(3):472–88, 2000.

[42] William J. Reeder. The benefits and costs of the section 8 existing housing program. *Journal of Public Economics*, 26(3):349 – 377, 1985.

[43] Ling Ren, Yan Zhang, and Jihong Solomon Zhao. The deterrent effect of the castle doctrine law on burglary in texas: A tale of outcomes in houston and dallas. *Crime & Delinquency*, 61(8):1127–1151, 2015.

[44] Donna Cohen Ross and Laura Cox. Beneath the surface: Barriers threaten to slow progress on expanding health coverage of children and families. Technical report, Kaiser Commission on Medicaid and the Uninsured, October 2004.

[45] Emmanuel Saez. Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2:180–212, August 2010.

[46] David Slusky and Donna Ginther. Did Medicaid expansion reduce medical divorce. Working Paper 23139, National Bureau of Economic Research, February 2017.

[47] Vernon K. Smith, Eileen Ellis, and Christina Chang. Eliminating the Medicaid asset test for families: A review of state experiences. Technical report, Kaiser Commission on Medicaid and the uninsured, April 2001.

[48] Sarah L. Taubman, Heidi L. Allen, Bill J. Wright, Katherine Baicker, and Amy N. Finkelstein. Medicaid increases emergency-department use: Evidence from oregon's health insurance experiment. *Science*, 343(6168):263–268, 2014.

[49] Daniel Webster, Cassandra Kercher Crifasi, and Jon S. Vernick. Effects of the repeal of missouri's handgun purchaser licensing law on homicides. *Journal of Urban Health*, 91:293–302, 2014.

[50] Aaron S. Yelowitz. The Medicaid notch, labor supply, and welfare participation: Evidence from eligibility expansions. *The Quarterly Journal of Economics*, 110(4):909–939, 1995.

[51] Aaron S. Yelowitz. Why did the SSI-disabled program grow so much? Disentangling the effect of Medicaid. *Journal of Health Economics*, 17(3):321–349, June 1998.

[52] Aaron S. Yelowitz. Will extending Medicaid to two-parent families encourage marriage? *The Journal of Human Resources*, 33(4):833–865, 1998b.