Essays in Public Economics

By

Manisha Jain

A dissertation submitted in partial fulfillment of

the requirements for the degree of

Doctor of Philosophy

(Economics)

at the

University of Wisconsin–Madison

2025

Date of final oral examination: 06/10/2025

The dissertation is approved by the following members of the Final Oral Committee:

Jeffrey Smith, Professor, Economics

Corina Mommaerts, Associate Professor, Economics

Laura Schechter, Professor, Economics

Priya Mukherjee, Assistant Professor, Agricultural and Applied Economics

Dedication

Dedicated to

my parents, Mridula Jain, and Madan Kishore Jain,

my sister, Madhavi Jain,

& Momo

Acknowledgments

First and foremost, I would like to thank my wonderful advisors: Corina Mommaerts, Jeff Smith, and Laura Schechter. They cared deeply about my papers, but more importantly, I always felt that they cared about me as a person.

I owe my deepest gratitude to Corina, whose mentorship has been a cornerstone of my academic journey. Her constant support and thoughtful engagement have shaped me as a scholar, instilling in me both rigor and curiosity. Through collaborative work and countless conversations, she has significantly (no t-test needed!) helped me grow into a better economist. I feel incredibly fortunate to have had her by my side throughout this journey. There is a well-deserved fandom around Jeff Smith, and it is easy to understand why. Jeff has been incredibly kind and patient, encouraging me to pursue ideas and to think from new perspectives, especially when I doubted whether something would work. Laura Schechter was the perfect addition to my committee, and I remain amazed by the time and effort she devotes to mentoring. She was instrumental in keeping me on track throughout this journey.

The second chapter of this dissertation would not have been possible without Corina Mommaerts and Jeff Weaver. I grew tremendously through this work, both in the process itself and from observing their inspiring work ethic and generous mentorship. Working with them has been not only great fun, but also an intellectually enriching experience that I will always treasure. Beyond this dissertation, I have had the privilege of collaborating with Richard Karlsson Linnér, Jessica Pac, and Christine Durrance. Even though those projects are not included here, working with them has helped me grow as a scholar, sharpen my thinking, and deepen my intellectual curiosity.

My gratitude also extends to Naoki Aizawa, Jesse Gregory, John Kennan, Rasmus Lentz, Priya Mukherjee, Martin O'Connell, Fernanda Rojas-Ampuero, Chris Taber, Kevin Thom, James Walker, and Matt Wiswall, who provided valuable feedback and encouragement at various stages of this process.

Many thanks as well to the department staff, especially Julie Anderson, Vicki Fugate, Becca George, Kim Grocholski, and Kelsey Hughes, for guiding me through the administrative side of the program over the years. A special thanks to Becca for her incredible support during the job market process.

I am deeply grateful to my friends and cohort members, including, but not limited to, Sossou Adjisse, Moshi Alam, Monica Agarwal, Priyadarshi Amar, Akarshik Banerjee, Shubhashrita Basu, Ashika Bhargav, Asha Bhardwaj, Saloni Bhogale, Matt Carl, Zhuoli Chen, Allie Greenman, Sakshi Gupta, Vikas Gawai, Karla Hernandez, Chunxiao Jing, Athipathi

Janarth, Melina Knabe, Utkarsh Kumar, Qin Li, Jaepil Lee, Liz Llanes, Nikitha Machineni, Julio Mereb, Rafeh Qureshi, Shreya Singh, Sakina Shibuya, Talita Silva, Dharanisree Srinivasan, Selidji Tossou, for their camaraderie, support, and countless conversations that made the past six years fun.

Finally, I owe everything to my parents, whose enormous sacrifices made it possible for me to come this far. And to my sister, who is always the first person I turn to whenever I face a problem, thank you for always being there for me.

Contents

List of Figures

List of Tables

Introduction

This dissertation is motivated by my interest in understanding socioeconomic disparities and is an attempt to make a small contribution to this literature. I study the impact of three policies that are designed to reduce socioeconomic disparities and emphasize how policy design and implementation are crucial to understanding their effects. Spanning topics in public economics, this collection of papers speaks to two overarching themes: (i) how response of fertility outcomes to cash transfers matters by policy design, and (ii) how targeted hiring subsidies and job quotas might not work as intended by policymakers.

Chapter 1 addresses the challenge of encouraging the birth and well-being of girls in smaller families, when son preference stems not only from perceived cost differences between boys and girls, but also from entrenched social norms favoring at least one son. I evaluate two conditional cash transfer programs in India that offer subsidies for girls in families with no more than two children, conditional on parental sterilization. One policy (the "girls-only" policy) restricts eligibility to families without sons; the other (the "girl-boy" policy) does not. I find that the girl-boy policy leads to an increase in subsidy-eligible child compositions, driven by a rise in one-boy-one-girl families. All-girl compositions do not increase under either policy, suggesting that the monetary incentives are insufficient to offset the perceived premium on having a son. Both policies increase sterilization at subsidy-eligible compositions, but the girls-only policy does so without altering fertility behavior, indicating that it may inefficiently incentivize a potentially risky medical procedure. These findings highlight the importance of designing fertility-related subsidies in a way that realistically accounts for the desire to have at least one son, particularly if increasing subsidy amounts is not feasible from a public expenditure perspective.

Chapters 2 and 3 shift the focus to labor market interventions for disadvantaged groups. Chapter 2, coauthored with Prof. Corina Mommaerts and Prof. Jeff Weaver, examines the Work Opportunity Tax Credit, a federal program in the US that covers up to 40% of wages for workers from targeted groups (such as SNAP beneficiaries) and subsidizes over two million hires annually. We find precise null effects on hiring, employment, and earnings across all our specifications, implying that the program functions largely as a transfer to (a small number of large) firms. We explore potential mechanisms and find that policy implementation plays a key role: although legally employers must establish a worker's eligibility at or before hiring, in practice, firms often silo this information from hiring managers, possibly to avoid legal risk associated with perceived discrimination against non-eligible candidates.

Chapter 3 returns to India to assess another redistributive intervention: public-sector job quotas for the Other Backward Classes (OBCs), a historically marginalized caste group. I find that the policy led to increased educational mobility among OBCs who were young enough to adjust their schooling decisions in response to the policy. Paradoxically, the job quotas decrease their occupational mobility. Thus, while the policy succeeded in raising educational attainment, likely due to raised educational aspirations among OBC youth, it did not translate into occupational mobility, even with increased employment in public-sector jobs.

Chapter 1

The Impact of Conditional Cash Transfers on Women's Fertility Choices and Girls' Outcomes

1.1 Abstract

India, like many other countries, is marked by persistent son preference and resulting gender inequalities. In response, policymakers instituted several Conditional Cash Transfer (CCT) programs to encourage the birth and education of girls in smaller families, along with sterilization. I evaluate two large CCTs that provide generous subsidies for girls, one that is restricted to families without boys and one without such a restriction. Using retrospective fertility data from the Demographic Health Survey and employing a difference-in-differences methodology, I find that the policy without the restriction to girls increases the probability of having a subsidy-eligible child composition by 8% relative to women in neighboring control states, driven by an increase in families with a composition that includes a boy. In contrast, all-girl compositions do not increase under either policy, suggesting that the demand elasticity for daughters is higher in the presence of a son. Both policies increase sterilization at subsidized compositions, but the girl-restricted policy does so without affecting fertility, suggesting that it inefficiently incentivizes women to undergo the risky procedure. However, the girl-restricted policy, by requiring subsidized girls to start school one year before the legally mandated age, increased early school enrollment.

1.2 Introduction

Many regions of the world, including countries in South Asia, Southeast Asia, the Middle East, and North Africa, are characterized by a persistent son preference, i.e., a parental bias towards sons rather than daughters (Dasgupta et al., 2003). For centuries, son preference has been associated with higher fertility rates, as parents often follow son-biased differential stopping rules, continuing to have children until their desired number of sons are born, which increases overall fertility (Arnold et al., 1998; Clark, 2000; Larsen et al., 1998). In recent decades, while desired family sizes have declined, son preference persists. The advent of prenatal sex-determination technology has enabled parents, particularly wealthier households who can access this technology, to maintain smaller families while still achieving their desired sex composition of children by selectively aborting female fetuses.[1] This creates a fertility–sex ratio tradeoff, where efforts to lower fertility rates may conflict with the goal of achieving gender balance. Son preference also results in higher post-natal discrimination against girls, manifesting in (1) higher female infant mortality due to neglect of their health and nutritional needs (Arnold, 1992) and (2) lower investment in their education (Pande and Astone, 2007).

Can conditional cash transfers (CCTs) influence fertility decisions, particularly regarding the sex composition of children and sterilization? Can the incentives they provide help address both pre- and post-natal discrimination against girls? Since the early 1990s, several Indian states have experimented with girl-promotion CCT programs, offering financial incentives to subsidize the birth and upbringing of girls in smaller families. These programs vary in design, with many limiting the number of boys, girls, and total children eligible for benefits, and requiring sterilization to ensure smaller family sizes. The motivation for these limits comes from the observed fertility-sex ratio tradeoff, even though these policies are targeted to poor households who are more likely to follow son-biased stopping rules than to have sex-selective abortions.[2] Another key feature of these policies is their long-term payment structure, linking transfers to milestones such as education and marriage, making it also important to assess whether they improve post-birth outcomes for girls, even if they

[1]Bhalotra and Cochrane (2010) estimate that approximately 0.48 million female fetuses were selectively aborted each year between 1995-2005 in India.

[2]Differential fertility strategy responses by household wealth are examined later in the paper.

do not alter reproductive behavior.

In this paper, I empirically examine the causal impact of subsidies provided under two such state-level programs, the Ladli Laxmi Yojana (LLY), which was started in the Indian state of Madhya Pradesh in 2007, and the Girl Child Protection Scheme (GCPS), which was started in the Indian state of Andhra Pradesh in 2005, on couple's reproductive decisions, specifically the sex composition of their children and sterilization choices, and on girls' outcomes, specifically, child mortality and enrollment in school. Both policies subsidize families with up to two children and require parental sterilization. These policies have a long-term payment structure, with more than 80% of the total payments made when the girl turns 18 or older, conditional on her being unmarried till that age and appearing in grade 12 examinations. The present discounted value of the subsidies offered per girl is substantial: about 9 times the average annual per capita expenditure for girls enrolled in LLY, and about 2.25-7 years of average annual per capita expenditure for girls enrolled in GCPS. These figures are notable when compared to the dowries that result in a cost differential between raising sons and daughters.[3]

There is a key difference in the eligible child compositions under these two policies: LLY allows eligibility even if one child is a boy (hereafter referred to as the 'girl-boy policy'), while GCPS restricts eligibility to families with only daughters (hereafter referred to as the 'girls-only policy'). Comparing these two policies is particularly important for studying fertility decisions because data from India suggest that most people have a revealed and stated preference for at least one son, meaning they generally do not prefer an all-daughter household (Jayachandran, 2015, 2017). In addition to a general aversion to daughters due to the financial burden of dowries, many families still hold a residual preference for having at least one son (Borker et al., 2024). This preference is often driven by cultural norms and expectations of old-age support from sons, who are traditionally seen as the ones to carry on the family name, retain family property, perform important religious rites (such as funeral rituals for parents in Hinduism) and provide care for aging parents (Drèze and Murthi, 2001; Pande and Astone, 2007).

[3]The average dowry payment in India ranges from one to two times the average annual household income (Chiplunkar and Weaver, 2023). Assuming hand-to-mouth consumption and an average household size of 4.8 in 2007, this translates to approximately 4.8 to 9.6 years of per capita expenditure.

Prior literature, including many economic theory papers (e.g., Ahn (1995); Becker (1960); Becker and Lewis (1973); Carro and Mira (2006); Heckman and Willis (1976); Hotz and Miller (1993); Li and Pantano (2023); Rosenzweig and Wolpin (1980); Willis (1973)), suggests that the demand for boys and girls depends on their cost, parental income, and parental preferences. Accordingly, subsidizing the cost of raising children based on sex and birth order has the potential to impact fertility behavior, but how effective these incentives are in practice would depend on their specific policy design. Thus, ultimately, it is an empirical question whether the current size of the subsidies under the present policy designs is sufficient to elicit a change in outcomes of women and girls. These programs, later replicated across multiple states, involve large public expenditures, yet their effectiveness has not been systematically evaluated.

I construct a balanced woman-year panel using retrospective birth history data from over 1.4 million women surveyed in the 2015-16 and 2019-20 rounds of India's Demographic Health Survey (DHS). The panel covers the years 2000 to 2014, aligning closely with the implementation of the policies studied. Additionally, I construct a girl-level dataset for girls born between 2000 to 2014, allowing me to track early childhood outcomes over time. I apply a difference-in-differences and event study methodology to the woman-year panel data and estimate the causal impact of these policies on women's fertility decisions, specifically, the likelihood of achieving a subsidy-eligible child composition, the total number of children, and the probability of sterilization at subsidized compositions. Similarly, I apply this empirical framework to the girls-level data to assess the effects on early childhood outcomes of subsidy-eligible girls, such as infant mortality and school enrollment.

I evaluate each policy separately, leveraging variation in incentives based on the year of program implementation and state of residence. Specifically, I compare women and subsidy-eligible girls in the treated state (i.e., the state where the policy was implemented) to their counterparts in a group of neighboring control states. I verify that the outcome variables trended in parallel prior to the policy change for the treated versus the control states. My main findings are as follows. First, the girl-boy policy increases the probability of women having a subsidy-eligible composition by 1.38pp, which is an 8% increase relative to the baseline of 17.21%, relative to women in neighboring states. This happens because of an increase in the proportion of women stopping at a one-boy-one-girl composition. Overall fertility declines, with the number of children decreasing by 0.04 children, i.e., a 3.6% decrease in comparison to a baseline mean of 1.16 children per woman. This effect is concentrated among younger women (aged 10-29 at the time of the policy implementation), poorer women, women with no prior children, and those with a prior boy. However, neither policy increases[4] the likelihood of families having an all-girls composition, reflecting a premium for the first son that is not compensated by the current subsidies. Second, both policies encourage sterilization at targeted compositions, but the girls-only policy does so without affecting fertility, making its payments inframarginal and inefficiently incentivizing women to undergo a risky procedure. Third, the girls-only policy promotes earlier school enrollment by requiring beneficiaries to start school a year before the legally mandated age (a 3.86pp, or a 5.9% increase relative to the baseline of 65% children being enrolled by age 5). This reflects how education-based conditionalities in transfer programs can influence the timing of school entry, though the broader effects on educational outcomes are yet to be explored.[5]

A number of robustness checks and falsification exercises support the causal interpretation of the policy effects. First, I demonstrate that the main fertility-related results are not driven by the choice of control states, by showing similar estimates using two strategies: (1) I restrict the analysis to bordering districts[6]. I include only districts in treated states that border control states, and compare them to the adjacent districts in the control states using a difference-in-differences framework. (2) I also apply the synthetic control method, constructing a "synthetic control" state by reweighting states in India that did not implement any CCT policy during the sample years, in order to match pre-exposure trends in the outcome variables. Next, I show that the policies do not affect the fertility behaviour of women aged 30-39 at the time of the policy implementation, who are nearing their completed fertility, or those whose existing child composition at the time of the policy implementation made them ineligible. Additionally, I provide evidence that the increase in sterilization at targeted compositions is not mirrored at non-targeted compositions, indicating that the

[4]The estimated effects are statistically and economically insignificant.

[5]The effects of early school enrollment on later-life outcomes present mixed evidence. While there is a huge literature on the importance of investing in early human capital formation (Cunha and Heckman, 2007), there is mixed evidence on how effective early school enrollment is in practice.

[6]Districts are geographical subdivisions of states in India, similar to counties in the US.

sterilization results are not driven by changes in other state-level sterilization policies.

I build a simple framework of fertility behavior that helps to rationalize and provide intuition behind the empirical findings. In the framework, women have preferences regarding the size of their family, the sex of their children, as well as sterilization, and I allow heterogeneity in the first two parameters. Women take two sequential decisions based on the number and sex composition of their existing children: first, whether to stop or continue childbearing, and second, whether to sterilize if they stop childbearing at her current child composition. I show that it is possible to have parameter values for different types of women in the population such that the policies might have only an observable sterilization effect at all-girls composition, but change fertility behavior when they already have a boy secured.

I contribute to several strands of literature. First, I add to the literature on son preference and missing women (Anderson and Ray, 2010; Dasgupta, 2005; Jayachandran, 2017, 2023; Jayachandran and Pande, 2017; Oster, 2005; Qian, 2008; Sen, 1992, 2003) by evaluating two policy designs motivated by the observation that fertility decline can exacerbate sex ratio imbalances at birth and lead to missing women. My findings emphasize the importance of targeting in policies aimed at encouraging the birth of girls, particularly in how the restriction of benefits to all-girls households influences women's fertility behavior. Two papers closely relate to my work. Anukriti (2018) examines a CCT program offering incentives to one-boy, one-girl, and two-girl couples,[7] conditional on sterilization. She finds that the probability of a woman having just one child increases by 0.6 to 1 percentage points, driven by an 11% increase in one-boy households, despite higher financial incentives for having a girl. This suggests that the demand elasticity for boys is greater than for girls. In contrast, the policy design in my study allows for a direct comparison between couples with and without a son, revealing that the demand elasticity for daughters is higher when a son is already present. The second closely related paper is Biswas et al. (2023), which examines a CCT program that provides benefits to all girls, without restricting the total number of children or boys. Their findings show that the probability of a girl's birth increases by 12% under the program,

[7]The present discounted value of the subsidy, expressed in terms of average annual per capita expenditure in the first year of implementation in the treated state, amounts to 14.94 years for parents with one girl, and 5.93 years for parents with one son or two daughters.

but this comes at the cost of higher fertility, as the probability of any birth increases by 3%.[8] In contrast, my findings suggest that allowing eligibility even with one boy can effectively promote girl births without increasing overall fertility.

More broadly, my findings contribute to the ongoing discussion of designing an optimal policy aimed at addressing both the birth and well-being of girls in smaller families (Anukriti, 2018; Bhartiya Janta Party, 2014). In contexts where son preference is driven not only by the cost differential between raising girls versus boys but also by deep-rooted social norms, policies targeting fertility decisions and outcomes of girls must be carefully calibrated. Additionally, the variation in Indian cultural and traditional values across states raises questions of external validity when applying similar policies elsewhere. Although Madhya Pradesh and Andhra Pradesh, the two states where the policies I study were implemented, differ in many respects, I present evidence that policy designs offering the current level of benefits solely to families with daughters may not be sufficient to influence fertility behavior in either state. Therefore, a key consideration in designing such policies is how "generous" they are in allowing families to exercise their preference for at least one son, particularly if increasing subsidy amounts is not feasible from a public expenditure perspective.

Second, my results speak to the broader literature on conditional cash transfer policies, especially those that study the impact of financial incentives on the reproductive behavior of women (Cohen et al., 2013; Desai and Tarozzi, 2011; González and Trommlerová, 2023; Kearney and Levine, 2009; Laroque and Salanié, 2014; Milligan, 2005; Walker, 1995). The CCT policies I look at provide financial incentives only to women with specific child compositions who undertake sterilization, rather than aiming to influence fertility and contraceptive behavior more generally. Another distinction is that the benefits in my context are not immediate but hinge on a promise of future payment, introducing a degree of risk for potential participants, as the success of these policies depends on long-term government funding and commitment. Third, I contribute to the literature examining the link between fertility policies and education outcomes, such as that explored by Qian (2009), in the context of China's one child policy, Pop-Eleches (2006) in the context of an abortion ban, and Ananat et al.

[8]Sinha and Yoong (2009) study another CCT scheme with a similar eligibility structure and find that the sex ratio becomes more gender-balanced, though fertility outcomes are not evaluated.

(2009) in the context of abortion legalization. A related literature looks at the intersection of education-based policies on fertility, such as Todd and Wolpin (2006), in the context of Mexico's Progresa policy. In this paper, I analyze the impact of an enrollment-based conditionality within what is primarily a fertility-related policy, focusing on the outcome of early school enrollment.

The rest of the paper proceeds as follows. In section 3, I present some facts about marriage and fertility in India, discuss potential reasons behind son preference, and discuss policies encouraging the birth of girls, with a focus on the two policies that I study. Section 4 provides a description of the data used in the analysis. Section 5 details the empirical strategy employed. Section 6 reports the estimation results and interprets the main findings. In Section 7, I build a simple theoretical framework to discuss the intuition behind the fertility-related outcomes of the two policies. Section 8 concludes.

1.3 Institutional Background

In this section, I provide a brief overview of marriage and fertility practices in India, discuss the underlying reasons behind son preference in more detail, and review the policies undertaken by the government to address the challenge of increasing the birth and well-being of girls in smaller families.

1.3.1 Marriage, Fertility, and Son-Preference Related Facts in India

I begin by presenting key facts and recent evidence on the nature of the marriage institution in India and fertility outcomes. Due to the social norm that all girls marry, marriage is nearly universal for females in India (Calvi et al., 2022). Bloch and Rao (2002) observe that "getting one's daughter married is considered an Indian parent's primary duty, and having an older unmarried daughter is regarded as a tremendous misfortune with significant social and economic costs". However, for the parents of the girl, marriage comes at a monetary cost - there is a transfer of wealth, called dowry, from the bride's family to the groom's family. Even though dowry has been illegal in India since 1961, dowry payments are almost ubiquitous and accepted as an entrenched institution, and it is generally difficult for a woman to marry without these transfers. Dowries remain remarkably substantial, with the average payment ranging from one to two times the average annual income of Indian households (Chiplunkar and Weaver, 2023). Marriage-related expenses are also disproportionately borne by the bride's family. The burden of dowry and marriage expenses is often cited as a critical reason for parents' aversion to raising daughters (Bhalotra et al., 2020), especially for credit constrained households. Conversely, sons are perceived as economic assets, as they bring dowries into the household.

Family-arranged marriages are still the prevailing method of matchmaking, with only 5%[9] of women reporting that they independently chose their spouse (Calvi et al., 2022). Marital dissolution, whether through legal divorce or informal separation, is exceedingly rare and often carries significant stigma. The divorce rate and the separation rate in the married population is just 0.24%[10] and 0.61% respectively (Jacob and Chattopadhyay, 2016). Remarriage is rarely exercised, with a pronounced gender differential: men are almost twice as likely as women to remarry, and there is a social stigma against remarriage for women, even widows. Additionally, having children outside of marriage is considered taboo, with only 0.61% of children born to unmarried women.[11] Marriages in India are typically patrilocal, with women leaving their paternal home to reside with their husbands and in-laws, often outside their natal village (Dyson and Moore, 1983). In the absence of formal social security institutions, this practice reinforces the expectation that sons, particularly the eldest, will provide financial support to their aging parents through coresidence. Daughters, on the other hand, are seen as *"paraya dhan"* (someone else's wealth), as once married, they no longer contribute to their natal family.

Figure 1.1a illustrates the probability of marriage, first birth, and any birth by age for ever-married women aged 15-49 at the time of the survey, using data from the 2015-16 round of the DHS. Indian women start marrying early, with most of this distribution,

---

[9] Data from the Indian Human Development Survey conducted in 2012.

[10] Data from the Census of India (2011) for the divorce and separation rate.

[11] Data from the 2015-16 and 2019-20 rounds of the DHS. This sample includes only ever-married women, but as marriage is nearly universal in India, the sample remains broadly representative of the population.

as shown by the green line, falling between the ages of 16-18.[12] The average and modal age at marriage for women is around 17.[13] Women typically begin childbearing soon after marriage due to cultural pressure. The probability of a first birth, shown by the orange line, reaches its mode around age 19, then declines, with very few first births occurring after age 30. Similarly, the blue line, which represents the probability of any birth (including both first and subsequent births), declines sharply after age 30,[14] highlighting that most women complete their childbearing in their twenties, with limited births occurring beyond age 30.

Couples generally aim to complete their families relatively quickly after the first child, with a median birth interval of 32 months. Although the use of modern temporary contraception remains limited, permanent sterilization becomes widespread once women have achieved their desired family size. By age 30, nearly one-third of women have undergone sterilization, as shown in Figure 1.1b. However, only about 1% of women report that their husbands have been sterilized, as men rarely opt for vasectomy due to the stigma surrounding the procedure, which is often perceived as emasculating (Merwin, 1974; Salve and Shekhar, 2023). India, the first country to launch a national family planning program in 1952, has historically relied on sterilization as a key component of its population control efforts, a practice that continues to play a central role in family planning today, despite the controversial mass forced sterilization campaigns during the 1970s under Prime Minister Indira Gandhi. In response to backlash from these campaigns, the government's approach shifted towards offering voluntary sterilization with monetary incentives, often targeted to poorer sections of the society. These one-time payments, offered immediately after the procedure and capped at around $22, display some variation across states. However, they do not vary based on the sex composition of the woman's children.

[12]The Prohibition of Child Marriage Act of 2006 sets the minimum legal marriage age at eighteen for women and twenty-one for men.

[13]The age at marriage for men is more variable, depending on the sex ratio in a region, but more than 80% of men marry by their thirties (Calvi et al., 2022).

[14]By age thirty, only 10% of women continue to give birth, and this probability drops further at older ages.

1.3.1.1 Son Preference and Fertility Outcomes

Son preference is associated with girls being disadvantaged and subjected to discrimination throughout their lives: foeticide, infanticide, gender gap in access to education (Pande and Astone, 2007), health and nutrition (Arnold, 1992), child marriage, and a lack of autonomy to make decisions. In this section, I present recent evidence on how son preference influences fertility outcomes.

Figure 1.2 shows the trends in the total fertility rate and the male-to-female child sex ratio (CSR) for children aged 0 to 6 years in India from 1960 to 2010. Over this period, the total fertility rate, depicted by the red line, declines steadily from above six children per woman in the 1960s to approximately 2.5 by 2010, reflecting broader demographic shifts toward smaller family sizes. In contrast, the CSR, represented by the blue line, exhibits an increasingly male-biased pattern, rising from around 105 male children per 100 female children in the early 1960s to 111 by 2010. This divergence begins in the late 1970s, coinciding with the legalization of abortion in 1971, the introduction of prenatal testing in the 1970s, and the widespread availability of ultrasound technology in the early 1980s. This trend demonstrates the fertility-sex ratio tradeoff mentioned in Section 1.2, where the goal of reducing fertility conflicts with maintaining a balanced sex ratio, a phenomenon also observed in other countries, such as in China.

Figure A1 displays the histogram of the number of children for women aged 30 and older[15] in 2015-16 in India. Approximately 30% of women have two children, 25% have three children, and only a small fraction (around 3%) have no children. In Figure 1.3, I look at the sex composition of the children born to these women using a heatmap. On the x-axis, I plot the number of girls and on the y-axis, I show the number of boys, with ”4” representing 4 or more children of each sex. Each cell reflects a different sex composition, with the percentage of women in that category displayed within the cell. For example, 3% of women have 0 boys and 0 girls, while 19% have 1 boy and 1 girl. One distinctive feature of Figure 1.3 is the clear imbalance in the sex composition of children, conditional on the total number of children. For example, let’s focus on women who have exactly two

[15] I consider these women to have completed fertility, based on the discussion related to Figure 1.1a.

children. The possible gender compositions for these women are two-boys-zero-girls (2B0G), one-boy-one-girl (1B1G), or zero-boys-two girls (0B2G).[16] If parents were indifferent to the sex of their children and concerned only with the total number of children, conditional on no differential gender mortality and having exactly two kids, we would anticipate about one-fourth of women with two children to have 2B0G, half to have 1B1G, and one-fourth to have 0B2G child compositions, as sex of a child at conception is naturally random, in the absence of sex-selective practices.[17] However, in Figure 1.3, conditional on having two kids, the probability of having a 0B2G composition is about 15%. Although this figure highlights these disparities, it does not definitively indicate whether parents achieve their preferred composition through son-biased stopping behavior or through selective abortion of female fetuses.

I investigate these theories by creating a woman-child level dataset based on the retrospective birth history for women aged 30+ in 2015-16 and analyzing two outcomes: the probability of bearing another child (parity progression) and the likelihood that the next birth is female. These are presented in Figures 1.4 and 1.5, respectively. Each figure contains four subpanels: the top-left shows results for the overall sample, followed clockwise by subgroups of women in the poorest, middle, and wealthiest wealth tertiles.[18] In Figure 1.4, I examine whether a woman's decision to continue childbearing depends on the sex composition of her existing children. The x-axis plots the number of children a woman has, disaggregated by whether they have at least one son or if they have no son. The y-axis plots the proportion of women who continue childbearing. Women with no children exhibit the highest probability of continuing childbearing, with nearly all continuing to have at least one child. As the number of children increases, the probability of having another child decreases overall. However, a clear son-preference emerges that aligns with the son-biased differential stopping rule: women with at least one son are more likely to stop childbearing, while women with only daughters are more likely to continue having children. For example, women with two daughters have around a 0.8 probability of having another child, whereas women with

[16]Here I am referring only to the number of children of each sex, not taking into account their birth order.

[17]1B1G can be obtained in two ways: BG, i.e., a boy followed by a girl, and GB, i.e., a girl followed by a boy.

[18]I construct a wealth index and divide women into three equal-sized groups: least wealthy, moderately wealthy, and most wealthy.

two children, atleast one of whom is a son, have a 0.55 probability of continuing childbearing. This disparity becomes more pronounced at higher parities, with the likelihood of continuing childbearing diverging sharply between those with no sons and those with at least one son. When disaggregated by wealth, important differences emerge. Across all sex compositions, poorer women are more likely to continue childbearing than wealthier women, which is consistent with higher completed fertility among the poor. Moreover, poorer women appear more likely to follow the son-biased stopping rule at any existing parity of children.

Next, in Figure 1.5, I explore the likelihood that the next child of these women will be female, conditional on the sex composition of their existing children. The x-axis mirrors that of the previous figure, but the y-axis now represents the proportion of women whose next child is female. In the absence of sex-selective abortions, we would expect the probability of the next child being female to be around 0.52 across all sex compositions. However, the data reveal a clear pattern. Women who already have at least one son tend to exhibit a probability closer to around 0.5 of having a daughter as their next child. But, women with only daughters, such as those with zero boys and two girls, show a markedly lower probability of having another daughter (0.44). These observed deviations from the natural 0.52 probability of conceiving a girl suggest that women are resorting to sex-selective abortions to ensure the birth of a son when they do not already have one. However, the likelihood that the next child is female remains closer to 0.5 across all sex compositions for poorer women, suggesting they are less likely to engage in selective abortions, perhaps, constrained by financial resources and lack of access. In contrast, wealthier families can afford sex-selective practices, enabling them to achieve their desired number of sons earlier in the fertility process. The policies I study aim to encourage the birth of girls in smaller families, but target poor households. Therefore, the impact of these policies is expected to work primarily by encouraging women to stop childbearing earlier, rather than discouraging sex-selective abortions of female fetuses.

1.3.2 Girl Promotion Policies

The central government and various state governments of India have instituted numerous policies aimed at encouraging the birth of girls, that I broadly categorize into (1) supply-side and (2) demand-side interventions. The central government implemented the Prenatal Diagnostic Techniques (Regulation and Prevention of Misuse) Act in 1994, a supply-side measure aimed at curbing sex-selective abortions by prohibiting foetal sex determination. However, the policy's impact was limited due to difficulties in enforcement,[19] as both abortions[20] and ultrasounds are legally permitted for reasons unrelated to sex selection.

Subsequently, policymakers shifted focus to demand-side policies designed to influence parental fertility behavior. These policies fall into two key approaches: (1) CCT programs that offer financial incentives to reduce the perceived cost of raising daughters, and (2) initiatives targeting the socio-cultural norm of son preference, which also underlies the lower perceived value of daughters. From a public expenditure perspective, CCT policies require greater financial outlays than supply-side measures or initiatives aimed at directly challenging son preference. However, CCT programs are expected to bring about more immediate behavioral changes, whereas shifts in long-standing social norms, such as son preference, tend to progress much more slowly. An example a of policy aimed at challenging the son preference norm directly is the "Beti Bachao Beti Padhao" (B3P) *(Save Daughters, Educate Daughters)* campaign, launched by the central government in 2015-16. B3P sought to raise awareness about the value of daughters through a comprehensive media campaign, including radio, television, print, and social media platforms. This effort was complemented by community outreach programs that include in-person door-to-door education efforts, mobile exhibition vans, field campaigns, and text messages seeking to reshape societal attitudes toward daughters.

[19] As of 2013, only 143 individuals had been prosecuted under this law.

[20] Abortion is legal under specific circumstances as defined by the Medical Termination of Pregnancy Act of 1971. These include cases such as risk to the mother's life, fetal abnormalities, rape, or contraceptive failure.

1.3.2.1 Girl-Promotion CCT Policies

Since the early 1990s, policymakers in India have experimented with girl-promotion CCT policies, with more than 20 such policies in existence, with varying designs.[21] The benefits provided by these programs are structured as either income support for the parents, or as subsidies for girls' education and marriage (which often functions as support for dowries). Ideally, in a context like India, where there is both a skewed male-to-female sex ratio at birth and a higher fertility than desired by policymakers (Ministry of Health and Family Welfare, 2000), the policy goal would be a reduction in the sex ratio at birth as well as the fertility rate. As a result, many of these policies are formulated to include a cap on the total number of children or sons that parents can have, for them (or their daughter) to be eligible for enrollment in these programs. This ensures that parents do not simply allow the birth of daughters, while continuing to pursue their desired number of children or sons. If the financial incentives are substantial enough, they may persuade parents who would otherwise opt for sex-selective abortions to retain their daughters, or to stop childbearing at smaller family sizes, even if they do not reach their desired number of sons. Many CCT programs also tie financial payments to compliance with post-birth conditions, such as birth registration, childhood immunization, school enrollment and retention, and delaying marriage of enrolled daughters until at least the legal age of 18 years. These conditions are added to address the gender-biased post-natal discrimination, and to ensure that the girls enrolled in the program are not simply neglected after birth.

While the broader goal of these policies is to promote the birth and well-being of girls, their design varies substantially in terms of the eligibility criteria and the size of financial incentives. Each policy is applicable only for parents residing in a particular state of India, and the eligibility for enrollment is limited to only the girls born after a specified date. Some programs, like the Girl Child Protection Scheme (GCPS) in Andhra Pradesh (AP) specifically target poor families, while others, such as the Dhanlaxmi scheme (implemented in 11 districts across 9 states), are open to all households regardless of their income. Apart from the date of birth of girls, the state of residence, and the income of parents, there

[21] Interested readers may refer to Anukriti (2018); Biswas et al. (2023); Sekher (2010, 2012) for details regarding policy design, subsidy amount, state and year of implementation, and eligibility information.

are three key dimensions along which eligibility in these policies varies: (1) the maximum number of girls per household who can enroll in the program, (2) the maximum number of children a family can have to remain eligible, and (3) the maximum number of sons allowed under the policy. Most programs restrict the number of girls per family who can receive benefits (e.g., two per family in GCPS, Ladli Laxmi Yojana (LLY), Mukhyamantri Ladli Laxmi Yojana (MLLY), Ladli; three per family in Bhagya Laxmi (BL)), though some, such as Dhanlaxmi, impose no such limits. When a policy imposes a limit on the maximum number of children, households are required to provide proof of sterilization of either parent and documentation of all children born before the procedure. However, some policies, such as Ladli and Dhanlaxmi, do not impose any restriction on the family size, and consequenty, do not require sterilization. Some programs, such as the GCPS, provide benefits only to parents who have daughters and no sons, while others, like the LLY, impose a limit of one boy, and the BL scheme imposes a limit of two boys.

In most of these schemes, with the exception of Devirupak, a substantial portion of the benefits (approximately 80-85%) is distributed when the girl reaches a certain age, typically between 18 and 21 years. This timing serves two key purposes. First, it ensures that the terminal conditions linked to the girl's welfare, such as completing high school education and avoiding child marriage, are met. Second, though not explicitly stated as a reason in these policies, the payments align with the period when dowry expenses typically occur. Thus, the CCT schemes are able to provide the credit-constrained parents of the girl with financial support during a time of significant financial need such as marriage, and highlight a way to reduce the aversion to the birth of a daughter stemming out of the future monetary burden associated with raising a girl child. For most policies, it is still too early to determine how beneficiaries are actually utilizing the terminal payments. However, a survey conducted by Sekher (2012) in Seoni district, Madhya Pradesh (MP), found that beneficiaries of the LLY believed the money was more likely to be spent on marriage and dowry than on higher education. Similarly, in states like Bihar and Jharkhand, the Dhanlakshmi scheme was predominantly noted as a way to support marriage expenses, with 90% of respondents expressing this view (Sekher and Ram, 2015). Consequently, social activists have critiqued these policies for covertly endorsing and perpetuating the institution of dowry (Sekher, 2010).

Smaller interim payments are linked to conditions like immunization or enrollment in certain grades in school. As a result, these policies represent a commitment to future payments, and households must trust the government to fulfill this promise in order to be incentivized to participate. Typically, the government invests the funds in the girl's name by issuing National Savings Certificates (NSCs), which mature after a specified period, providing proof of the commitment. However, families are not allowed to take out loans or credit against these certificates, and early payments are not permitted under any circumstances. Most of these schemes are administered by the Department of Women and Child Development, leveraging the vast network of Integrated Child Development Services (ICDS) and *Anganwadi* (community health workers). *Anganwadi* workers play a key role in motivating parents to register for the schemes; e.g., in MP's Seoni district, 80% of scheme enrollments were facilitated by *Anganwadi* workers (Sekher, 2012). In some cases, such as the LLY in MP, local *panchayat* (village council) leaders have also contributed by promoting the schemes (for example, through organizing special *Gram Sabha* meetings[22] to raise awareness) and identifying potential beneficiaries. Many state governments take pride in implementing these schemes, often presenting them as major achievements. The success of the LLY in MP (in terms of enrollment numbers) prompted the Bharatiya Janata Party (BJP), one of the major political parties in India, to propose a similar scheme in its manifesto during the 2014 parliamentary election, with a promise to implement it nationwide if they came to power.[23]

I examine two policies in detail to assess the nuances of policy design: the Ladli Laxmi Yojana (LLY) in Madhya Pradesh and the Girl Child Protection Scheme (GCPS) in Andhra Pradesh. I select these two schemes for the following two reasons. First, both policies aim to promote the birth and welfare of girls, but the design differs significantly in terms of eligible family composition. In both GCPS and LLY, girls can be enrolled only if the family size does not exceed two children. However, GCPS is stricter; if a boy is born, the family loses eligibility for all subsidies. In contrast, LLY allows for the birth of a boy without forfeiting the benefits for the girl, although the boy does not receive any subsidy. This makes LLY more flexible, while GCPS is focused on incentivizing families with only girls. Second, the

[22]General assembly meeting of all registered voters in a village.

[23]The BJP won the 2014 election and formed the central government, but no such nationwide scheme was subsequently implemented.

timing of these policies is well-aligned with the available data. While similar policies, such as the GCPS scheme in Tamil Nadu in the 1990s, were introduced earlier, using data from that period would require older survey rounds with smaller sample sizes, potentially leading to underpowered analyses. On the other hand, policies modeled after LLY, such as MLLY in Jharkhand were introduced later, leaving fewer post-treatment periods for analysis. This limits the ability to fully capture their long-term effects, as it is likely that the policies take time to exhibit measurable impacts.[24] I hereafter refer to the LLY policy as the girl-boy policy and the GCPS as the girls-only policy.

Girl-Boy Policy (LLY) The girl-boy policy that I study, called the Ladli Laxmi Yojana (LLY), is a CCT scheme instituted by the government of Madhya Pradesh (MP) to simultaneously tackle the problems of high fertility and high sex ratio at birth. The program was announced and went into effect in all the districts of the state on April 1, 2007. Girls born after January 1, 2006 to couples with at most one prior surviving child can be enrolled under LLY. Parents can thus enroll a maximum of two daughters[25] in this scheme. Enrollment of a girl is allowed up to one year after the last child birth of a woman upon the production of a family planning (sterilization) certificate, birth certificate of the girl, state-of-domicile certificate, photograph of the girl with her mother, and an undertaking that her parents are not income tax payers.[26] The government of MP issues National Saving Certificates (NSCs) worth Rs. 6,000($150) each for five consecutive years in the name of the enrolled girl. Thus, a total of Rs. 30,000($750) is invested on her behalf. The payment structure of LLY is listed in Table 1.1. Altogether, the incentives total Rs. 118,300($2957), and assuming a discount rate of 3%, the present discounted value of these benefits is Rs. 65,300($1632). The present discounted value of these benefits amounts to 8.58 times the average annual rural per capita

[24]Mechanically, it takes time for households to have children and fulfill the eligibility conditions of the policy.

[25]Technically, a couple can enroll three girls under LLY as well. This is possible under the following two scenarios. If a woman who has no prior children gives birth to triplet girls, then all the three girls would receive LLY's benefits, as long as the other eligibility criteria are met. Similarly, twin girls born to a woman with one prior child can also be registered under LLY, irrespective of the sex of the previous child.

[26]Given India's low per capita income and high tax exemption thresholds, the vast majority of households, except those in the top 1 percent of the income distribution, are exempt from personal income tax (Banerjee and Piketty, 2005). As I do not observe households' income-tax status in the data, I examine differential effects of the policy by wealth groups, under the assumption that the subsidy is more meaningful to less wealthy individuals.

expenditure in Madhya Pradesh in 2007.

Girls-Only Policy (GCPS) The girls-only policy that I study, called the GCPS, was launched by the government of Andhra Pradesh (AP)[27] on March 8, 2005. The scheme is operational across the state, excluding two blocks where the Dhanlakshmi[28] scheme is in effect. Eligibility for the GCPS is restricted to families with either a single girl child or two girl children, provided one of the parents has undergone sterilization. Furthermore, the family's annual income must be below Rs. 20,000($400) in rural areas and below Rs. 24,000($480) in urban areas.[29] Applications must be submitted before the second girl turns three years old, or before the only daughter turns three in the case of single-child families. Table 1.2 summarizes the conditions, stages and installments of the payments provided under this policy. Consequent to enrollment, enrolled girls must be admitted to school by age five. Compared to the average annual rural per capita expenditure in Andhra Pradesh in 2005, the present discounted value of these benefits equates to 6.69 years of such expenditure if one girl is enrolled, and 4.5 years of such expenditure if two girls are enrolled.[30]

1.4 Data

The data used in the main analysis of this study come from the 2015-16 and 2018-19 rounds of the Demographic Health Survey (DHS) of India, which is a nationwide, repeated cross-sectional survey, representative at the state level. The survey encompasses questions on a wide range of socio-economic and demographic characteristics, for example, each household member's age, residential status, education, caste, religion, wealth, employment, expenditure, etc.[31] Ever-married women between the ages of 15 and 49 years are asked additional

[27]Telangana was created from AP in June 2014 and decided to continue with the policy. Since my sample ends in 2014, I refer to the treated state as AP, though technically, two states were treated after 2014.

[28]Launched in 2008, Dhanlaxmi was implemented in 11 blocks across India with the most adverse SRBs. This program provided benefits, which are comparable in magnitude to those under LLY, to all girls born between 2008 and 2013, without imposing restrictions on the number of boys or total children in a family.

[29]These limits were revised in 2013 to Rs. 40,000($800) and Rs. 48,000($960), respectively, for children born on or after January 3, 2013.

[30]The total subsidy amount is smaller when two girls are enrolled than when only one girl is enrolled: each girl receives the equivalent of 2.25 years of expenditure, leading to a combined total of 4.5 years.

[31]Residential status, education, caste, religion, wealth, employment, and expenditure are measured at the time the survey is conducted.

questions related to health, marriage, fertility, and gender relations. Of particular importance is the birth history of the interviewed women, complete to the year of the survey, i.e., the birth order, sex of the child, year of birth, woman's age at each child birth, and the child's age at death (if the child is deceased). Women are also asked about their sterilization status, including the year and the month of sterilization, if applicable. Information on usage of other contraception methods at the time of the survey is also reported. Additional information is available for recent births, such as immunization status.

There are two advantages of using data sourced from DHS. First, over 1.4 million women were surveyed across these two rounds. The second round of the Indian Human Development Survey (IHDS-2) also contains all the socio-economic and demographic characteristics described above, but the number of interviewed woman is much smaller, about 40,000. Second, the data record the complete birth histories of women, as opposed to datasets like the District Level Household and Facility Survey (DLHS) and Annual Health Survey (AHS). The later rounds of DLHS do not report the complete birth histories of interviewed women, but rather collect information for only recent births.[32] However, the DHS data, as any other publicly-available household survey data in India, does not explicitly include actual participation in these policies, and thus, I am limited to estimating an intent-to-treat effect, relying on the statutory eligibility criteria for the policies.

Following the procedure outlined in Anukriti (2018), I use the retrospective birth histories from the two cross-sections of interviewed women to construct a woman-year panel. However, unlike Anukriti (2018), I do not condition the entry of women into the sample on their year of marriage, but rather construct a balanced sample, where each woman exists in the panel regardless of whether she was married by that age. Age at marriage might be potentially affected by these policies, as women might marry and have children earlier, if they anticipate the policies might be discontinued in the future. Creating a balanced sample also prevents any compositional changes in sociodemographic characteristics that can influence fertility behavior of women. Since both policies applied only to permanent residents of the implementing state, I exclude all women identified in the survey as visitors,

[32]DLHS-3 and DLHS-4 (conducted in 2007-08 and 2012-13 respectively) recorded information only for births after January 1, 2004 and January 1, 2008 respectively.

i.e., those who were staying in the household the night before the interview but were not usual members of the household. My panel comprises the years from 2000 to 2014. I select 2000 as the starting year in the panel to ensure that the period under consideration is closer to the implementation of the policies, thereby increasing the likelihood that the outcomes being examined occurred in a context where other influencing factors, like the availability of pre-natal sex determination, remained relatively similar. I limit my panel to the years prior to 2015 to maintain a balanced panel[33] and to avoid the influence of the *Beti Bachao Beti Padhao* (B3P) program, that was launched by the Central Government of India in 2015. B3P seeks to educate households about the equal value of girls and boys, and could therefore influence fertility decisions of women or the post-birth investment and outcomes of girls. I remove women who were younger than 10 years old in the year the policy was implemented, as they are highly unlikely to have children in the pre-policy years.

Additionally, I create a girl-level dataset for the children of women in this panel, tracking early childhood outcomes, for tracking early childhood outcomes for each girl in the 2000–2014 birth cohort. I then restrict my sample to policy-eligible girls, where I define policy-eligibility based on the number of siblings at birth: i.e. girls who had at most one brother and one sister for the girls-only policy and girls who had at most one sister at the time she was born for the girls-only policy. For each policy-eligible girl, I observe post-birth early childhood outcomes, specifically, whether they (1) survived to age one, (2) survived to age five,[34] and (3) enrolled in school by age five.

1.4.1 Treatment Years and Control States

I now discuss the choice of treatment years and control states. As will be elaborated later in Section 1.5, I use a difference-in-differences framework to estimate the causal impact of the girl-boy and the girls-only policies on the fertility decisions and post-birth outcomes of women and their children. Since both policies were implemented simultaneously across all districts in their respective states, I am unable to leverage within-state geographical variation in program

[33]Although DHS-4 interviewed women in 2015-16, some states conducted interviews early in 2015, which means that I would have their fertility history only till the end of 2014.

[34]I include only girls for whom I can observe their status five years post birth. Since I am using two rounds of the DHS survey, one of which was conducted in 2015, I exclude information from this survey for girls who were born after 2010 so that my data are not truncated.

implementation for my empirical strategy. Instead, I use nearby states as control states, as they are likely to be similar to the treatment state in terms of patriarchal institutions and socio-cultural norms, such as kinship structures, marriage customs, patrilineality, female autonomy, and the organization of the agrarian economy. This institutional setup influences the cultural son-preference norm, and is, consequently, expected to affect fertility decisions of women and post-birth outcomes of girls. Additionally, I exclude districts and states that had a girl-promotion CCT policy active in any of the years from 2000-2015. Finally, I restrict my analysis to neighboring states that minimize differential pre-trends in the main outcome variables. This is crucial for satisfying the key identifying assumption of the difference-in-differences method, which requires parallel trends to credibly interpret the results.

After applying these three criteria (details in Section A.3.1), I obtain Bihar, Chhattisgarh, Gujarat, Odisha, Rajasthan, and Uttar Pradesh as control states for the girl-boy policy (Figure 1.6(a)) and Chhattisgarh, Gujarat, Kerala, Maharashtra, and Odisha as the control states for the girls-only policy (Figure 1.6(b)). I exclude the districts listed in Table A1 in these states, as they were part of another girl-promotion CCT scheme called Dhanlaxmi. Launched in 2008, Dhanlaxmi was implemented in the 11 blocks[35] across India with the most adverse sex ratio at birth. This program provided benefits to all girls born between 2008 and 2013, without imposing restrictions on the number of boys or total children in a family, which could confound the effects of the policies I study.

Even though the girl–boy policy was announced in April 2007, I treat 2008 as the first treatment year for fertility-related outcomes. The observable fertility outcomes are subsequent births and the sex of children born. Children conceived after the policy announcement would be born in 2008, and female births in 2007 were unlikely to have been influenced by the policy, as late-term sex-selective abortions carry high risks. Similarly, although the girls-only policy was announced in 2005, I take 2006 as the first treatment year for fertility outcomes. For analyses focusing on girl-specific outcomes, I use the first treatment year based on the policy's eligibility rule rather than the year of announcement. Under the girl–boy policy, all girls born in or after 2006 are eligible, so I use 2006 as the first treatment year. For the girls-only policy, which covers girls born in or after 2005, I use 2005 as the first treatment

[35] A block is a subdivision of a district.

year.

1.4.2 Descriptive Statistics

Tables 1.3 and 1.4 present the sample means of key variables, separately for treatment and control states in the pre-policy years, for the girl-boy policy and the girls-only policy, respectively, using the woman-year panel after applying the appropriate age restriction for women as outlined in the previous section. To study the effect of the policies on women's fertility decisions, I first construct an indicator variable that takes the value of one if the woman's child composition at the end of year t is subsidy-eligible. For the girl-boy policy, subsidy-eligible compositions mean one-boy-one-girl (1B1G), zero-boys-one-girl (0B1G), or zero-boys-two-girls (0B2G). For the girls-only policy, eligibility is limited to compositions with zero-boys-one-girl (0B1G) or zero-boys-two-girls (0B2G), excluding 1B1G. My other main variable is sterilization at subsidy-eligible compositions, which takes the value of one if the woman is sterilized at a subsidy-eligible composition at the end of year t. I also present summary statistics for other variables that could potentially influence women's fertility decisions.

Even though I find statistically significant differences between the treated and the control states in several variables listed for both policies, these differences do not pose a direct threat to identification in the difference-in-differences strategy that I use later. The key identification assumption is not identical baseline characteristics; the identification assumption is that in the absence of the policy, the average outcomes of women in the treatment state would follow parallel trends to those in control states, conditional on the included covariates. However, I include the socio-demographic characteristics (urban/rural status, education, religion, caste, wealth index) as controls in my regressions to ensure that the effects of the policies are not confounded by differences in these variables between the treated and control states. Since the sample is balanced by construction, there should not be compositional changes in these variables over time, as these socio-demographic characteristics are measured at the time the survey is conducted.[36]

[36]These characteristics are unlikely to change over the woman's life. For example, most women do not acquire additional years of schooling post-marriage. Wealth index, when measured in categories in a relative

1.5 Empirical Strategy

The objective of this study is to empirically test whether the financial incentives provided under the girl-boy and girls-only policy have any causal impact on women's fertility decisions and the post-birth outcomes of policy-eligible girls. I employ a difference-in-differences and event study methodology to the woman-year panel data and estimate the causal impact of these policies on women's fertility decisions—specifically, achieving a subsidy-eligible child composition, the total number of children, and sterilization at subsidized compositions. Similarly, I apply this empirical framework to the girls-level data to assess the effects on early childhood outcomes of subsidy-eligible girls, such as infant mortality and school enrollment. I evaluate each policy separately, leveraging variation in incentives based on the year of program implementation and state of residence. Specifically, I compare the temporal variation in the outcomes of women and subsidy-eligible girls in the treated state (i.e., the state where the policy was implemented) to their counterparts in a group of neighboring control states. I define policy-eligibility for girls based on the family composition at the time of their birth. For the girl-boy policy, a girl is eligible if she has at most one sibling at birth, regardless of the sibling's sex. For the girls-only policy, a girl is eligible if she has at most one sister and no brothers at the time of her birth. Even though this strategy is not entirely clean, since the child-outcome dataset is restricted to girls who were policy-eligible at birth and continued eligibility at the time outcomes are measured may depend on fertility decisions influenced by the policy, it remains a worthwhile exercise to undertake.

Formally expressed, I estimate the following two-way fixed effects (TWFE) difference-in-differences equation:

$$y_{ist} = \beta[Treat_s \times 1\{t \geq \tau\}] + \boldsymbol{X_i'}\delta + \gamma_s + \theta_t + \nu_{a(\tau)} + \varepsilon_{ist} \tag{1.1}$$

Here, y_{ist} is the outcome of interest of woman i at the end of year t who resides in states s, or the outcome of interest for girl i born in year t residing in state s. $Treat_s$ is an indicator for whether the state in which she resides is a state in which the CCT

sense, is also unlikely to change.

scheme was implemented; τ is the first treatment year; $\boldsymbol{X_i'}$ is a vector of covariates: woman's education (or mother's education for girl outcomes), indicators for sex of the household head, residence in an urban area, household's religion (Hindu, Muslim, Sikh, Christian, Others), caste (General, Scheduled Caste (SC), Scheduled Tribe (ST), Other Backward Castes (OBC)), and wealth index (1 to 5, with 1 being the lowest)); γ_s are state fixed effects; θ_t are year (or year of birth for the girl regressions) fixed effects; $\nu_{a(\tau)}$ are age-at-the-policy fixed effects (i.e. birth cohort-time fixed effects) (only included in the woman regressions); and ε_{it} is the error term.

The coefficient of interest in Equation (1.1) is β, i.e., the coefficient on the interaction term $Treat_s \times 1\{t \geq \tau\}$, which captures the differential trend in the outcome variable between the treatment and the control groups. The identification of β hinges on the assumption that, in the absence of the policy, the average outcomes of women (or policy-eligible girls) in the treatment state would follow parallel trends to those in control states, conditional on the included covariates. All my regression estimates provide intent-to-treat effects as I do not have information on actual program participation.

I also estimate the event study counterpart of Equation (1.1), which allows me to visually (and flexibly) examine the differential patterns in the outcomes between the treated and controls, relative to the time of policy implementation:

$$y_{ist} = \sum_{z=\tau-2000, z\neq -1}^{2014-\tau} \beta_z Treat_s \times [1\{Year_t - \tau = z\}] + \boldsymbol{X_i'}\delta + \gamma_s + \theta_t + \nu_{a(\tau)} + \varepsilon_{ist} \qquad (1.2)$$

Here, 2000 and 2014 demark the range of the years in the sample. The coefficients of interest are now the β_z's , which represent the effect of the policy at various time points z years after its implementation, compared to the baseline year of -1 (the year immediately prior to the policy's introduction). These coefficients capture the dynamic impacts of the policy over time relative to the baseline period, enabling a clearer understanding of the policy's effects as they evolve after implementation. Given the small number of treated clusters (one treated state per policy evaluation), conventional inference (based on p-values from asymptotic distributions) may fail, thus, I use wild bootstrap standard errors, clustered at the state level (MacKinnon et al., 2023) in Equations (1.1) and (1.2).

I conclude this section with an important caveat regarding the assignment of treatment status. In this analysis, I have inferred a woman's (or girl's) treatment status based solely on her state of residence at the time of the survey and the year in question. However, as discussed in Section 1.3.2, eligibility for the girl-boy policy is conditional on the parents not being income taxpayers, while beneficiaries of the girls-only policy must belong to lower-income households. Therefore, it is likely that the fertility decisions influenced by the girl-boy policy primarily affect non-taxpaying households, and those impacted by the girls-only policy are concentrated among the poorer segments of society. Unfortunately, I do not have data on household income or tax status, which would have allowed for a more precise identification of eligible families, and this is a limitation of the present study, though I do have data on their wealth reported at the time of the survey. Rather than excluding women based on an arbitrary wealth threshold, I estimate Equations (1.1) and (1.2) for the entire sample of women. I later explore heterogeneity by conducting separate analyses based on wealth quintiles.

1.5.1 Border Strategy

As a robustness check, I narrow the analysis to districts in the treatment state that border neighboring states, using the adjacent districts in those neighboring states as the control group. This approach leverages the fact that bordering districts are more likely to share similar sociocultural norms, making them more comparable. To avoid contamination, as in the main analysis, I exclude any bordering districts that implemented a different girl-promotion CCT policy, ensuring that the measured impact is attributable solely to the two policies of interest. I then apply the difference-in-differences method to this restricted sample to assess whether the results are consistent with those from the main empirical strategy. Details of the bordering districts are presented in Appendix A.3.2.

1.5.2 Synthetic Controls

I use the synthetic control method (Abadie et al., 2010; Abadie and Gardeazabal, 2003) as an additional robustness check to verify that the results are not sensitive to the choice of the control states. The idea here is to use a larger pool of donor states, which are all states in India that did not have a similar CCT promotion scheme during 2000-2014, to create a "synthetic" control state by reweighting the donor states so as to match the pre-exposure trends in the outcome variable and covariates. These donor states (all states in India without a girl -promotion CCT policy) are Jammu and Kashmir, Rajasthan, Uttar Pradesh, Chattisgarh, Gujarat, Maharashtra, Kerela, Orissa, Bihar, Sikkim, Arunachal Pradesh, Tripura, Manipur, Meghalaya, Mizoram, and Nagaland. To implement the synthetic control method, I collapse my individual level data to the state-year level. For instance, the dummy variable for a subsidy-eligible composition is now the proportion of women in state s in year t with a subsidy-eligible child composition. Similarly, a covariate like a dummy for caste affiliation now becomes the percent of women who belong to that caste group in that state-year.

1.6 Results

I organize my results in two subsections: in Section 1.6.1, I present the effect of the policies on fertility-related outcomes of women, specifically, their child composition, number of living children, and sterilization status, using data from the woman-year panel. In Section 1.6.2, I present the impact of the policies on the post-birth early-childhood outcomes of policy-eligible girls, specifically, infant mortality rate, under-5 mortality rate, and school enrollment, using data from the child-level dataset.

1.6.1 Effect on Fertility Related Outcomes

1.6.1.1 Targeted Child Compositions

To assess the behavioral responses in reproductive behavior, I begin by examining the probability that a woman achieves a subsidy-eligible, or, in other words, a targeted child composition by the end of year t, as specified by the policy. Recall that for the girl-boy policy, subsidy-eligible compositions mean one-boy-one-girl (1B1G), zero-boys-one-girl (0B1G), or zero-boys-two-girls (0B2G). For the girls-only policy, eligibility is limited to compositions with zero-boys-one-girl (0B1G) or zero-boys-two-girls (0B2G). Figure 1.7 presents the main results on the effect of the policy on targeted child compositions. The green line in Figure 1.7 illustrates the event study for the girl-boy policy, derived from Equation (1.2), and demonstrates no evidence of pre-trends in the outcome variable, thereby supporting the parallel trends assumption. Similarly, the blue line in Figure 1.7 presents the event study for the girls-only policy, also based on Equation (1.2), showing that the treated state and its control states followed comparable trajectories prior to the implementation of the girls-only policy. Table 1.5 reports analogous difference-in-differences estimates, with the first column reflecting the impact of the girl-boy policy and the second column presenting the effects of the girls-only policy. The results indicate that the girl-boy policy increases targeted child compositions by 1.38 percentage points, representing an 8.01% increase, relative to the mean of the dependent variable in the control states before the policy implementation (17.21%). In contrast, the girls-only policy yields no statistically significant effect, with a precise null estimate allowing me to reject effect sizes greater than 0.4 percentage points. I also observe in Figure 1.7 that the effect of the girl-boy policy on targeted child compositions increases steadily over time. This is expected, given the time required to plan and have children, particularly for households starting with no children who need to reach a 0B2G or 1B1G composition. Additionally, the growing impact may result from information diffusion, as households gradually become aware of the policy. This may occur through women observing people they know applying for the program, or as awareness spreads within the community via increased government promotion or advertisement of the policy.

I verify that the estimates for both the policies are of the same magnitude as in Columns 1 and 5 (i.e., those labelled 'Overall') of Table A2 when using the border strategy, i.e., by restricting the analysis to bordering districts of the treated state. I also verify that the results look similar when using the synthetic control method, as evidenced in Figure A3.

Next, I investigate which specific child composition(s) drive the increase in the percentage of women with targeted child compositions under the girl-boy policy. The outcome variables in this analysis are four separate binary indicators, one for each child composition:

0B1G, 0B2G, 1B1G, and Rest.[37] Each variable takes the value of 1 if a woman achieves the corresponding composition by the end of year t, and 0 otherwise. Figure 3.4a reports event study estimates from Equation (1.2) for each of these four outcome variables. The pink line corresponds to the binary indicator for the 0B1G composition at the end of the year t, the magenta line represents the 0B2G composition, the purple line corresponds to 1B1G, and the black line represents the rest of the compositions. For all four outcomes, there is no evidence of differential pre-trends between the treated and control states prior to the implementation of the girl-boy policy. Figure 3.4a clearly shows that the increase in the likelihood of women attaining a targeted child composition under the girl-boy policy is driven solely by a rise in the 1B1G composition (a large and statistically significant increase in the proportion of women of 1.55pp). There are no large effects on the likelihood of having a 0B1G (in fact, Column (1) of Table 1.6 shows that this probability decreases by 0.10pp) or 0B2G composition in the treated state (an increase of 0.06 pp as in Column (3) of Table 1.6) relative to the control states, after the policy change. These coefficients are tightly estimated and rule out economically meaningful effects. On the other hand, the black line shows a robust decline in the likelihood of women having non-targeted compositions in the treated state compared to the control states. This indicates that households are more inclined to have a one-boy-one-girl composition, substituting away from non-targeted compositions.

Heterogeneity based on Eligible Women for Girl-Boy policy The observed increase in the percentage of women with a 1B1G composition after the implementation of the girl-boy policy is interesting, but we should be cautious at this stage in interpreting this as a causal effect of the policy alone, since the birth order is not accounted for in this outcome. Specifically, this outcome simply captures whether a woman has one boy and one girl, regardless of the order in which they were born. Thus, a woman with a girl born before the policy and a boy born afterward is treated the same in this analysis as a woman with a boy born before and a girl born after, even though only the latter would qualify for policy benefits, as the policy conditions specify that only girls born after the policy are eligible for benefits. Therefore, it is crucial to explore the mechanisms behind the observed increase in

[37]Specifically, "Rest" refers to child compositions not targeted by the girl-boy policy, i.e. no children, B, BB, or more than two children at the end of year t.

the proportion of women with a one-boy-one-girl composition after the introduction of the girl-boy policy, based on the woman's child composition at the start of the policy.[38] Three types of birth histories at the policy's start could make subsequent girl births eligible for benefits: no children, one boy, or one girl. Women without children could potentially earn benefits for having two daughters, while women with one child, whether a boy or a girl, would only be eligible for benefits if their next child was a girl. Additionally, to qualify for the policy, families must also stop at a maximum of two children, as the benefits are contingent upon limiting family size to no more than two children. Thus, I divide my sample into four categories based on the woman's eligibility at the start of the policy: women with no children, women with one boy, women with one girl, and women with two or more children. This breakdown enables me to examine how different groups responded to the girl-boy policy and helps identify which marginal groups were most affected.

First, in Figure 3.4b, I present event study estimates for the sample of women who had no children at the start of the policy. Each line represents a separate binary outcome variable at the end of year t: green for OB0G (zero boys and zero girls, i.e., childless), light blue for 1B0G (one boy and zero girls), dark blue for 2B0G (two boys and zero girls), pink for 0B1G (zero boys and one girl), magenta for 0B2G (zero boys and two girls), purple for 1B1G (one boy and one girl), and black for more than two children. Unsurprisingly, there are no differential pre-trends, which can mostly be explained by the fact that women with zero children at the time of the policy also had zero children in the years leading up to the policy, resulting in zeros across all outcome variables except for the 0B0G outcome.[39] The girl-boy policy reduces the expected cost of having children, which in turn lowers the probability of women in the treated state remaining childless relative to control states. While there are no statistically or economically significant changes in behavior regarding having exactly 1B0G, 2B0G, 0B1G, or 0B2G compositions, the policy leads to an increase in 1B1G compositions and a decrease in the likelihood of having more than two children. This indicates that women

[38] Even though girls enrolled in the girl-boy policy could be born in or after January 2006, I use 2007 to define eligibility, as decisions made before the policy announcement in April 2007 would not have been influenced by it.

[39] Note that child composition is based on the number and sex composition of living children, so it is possible for a woman with zero children at the policy start to have given birth to children before, but they may have passed away by the time the policy began.

are more likely to shift toward a 1B1G composition, both by the girl-boy policy encouraging women to have children rather than remain childless[40] and by stopping at 1B1G rather than moving to higher parity.

Second, I examine the effect of the girl-boy policy on women who had exactly one boy and no girls at the start of the policy. Figure 3.4c presents event study estimates for this sample, with each line representing a separate binary outcome variable at the end of year t: light blue for 1B0G (one boy and no girls), dark blue for 2B0G (two boys and no girls), purple for 1B1G (one boy and one girl), and black for more than two children. The results show that the policy increases the likelihood of a woman attaining a 1B1G composition. Given that these women already had a boy before the policy, this suggests that the girl in this composition is eligible to be enrolled in the policy. There is a concurrent decrease in the probability of women having more than two children, indicating that women in this group are opting to stop at the 1B1G composition rather than progressing to higher parity. This suggests that the girl-boy policy is successful in encouraging families to limit their fertility once they attain the one-boy-one-girl composition.

Third, I assess the impact of the girl-boy policy on women who had one girl and no boys at the start of the policy. Figure 3.4d illustrates the event study estimates for this group, with pink representing the 0B1G composition (no boys, one girl), magenta representing 0B2G (no boys, two girls), purple for 1B1G (one boy, one girl), and black for families with more than two children. Despite the potential incentive for these women to have a second girl (since a girl born post-policy would be eligible for benefits) the results indicate no notable changes across all four child composition outcome variables; women with one girl before the policy in the treated state do not display any statistically significant or large shifts in their fertility patterns compared to women in the control states, either before or after the policy, suggesting that the policy had no influence on their childbearing strategy. Taken together, these results suggest that the observed increase in the 1B1G composition under the girl-boy policy is not driven by women who already had a girl before the policy and subsequently

[40]It is important to note that not all women in my data have completed fertility, as the sample includes younger women. This shift could simply reflect the policy encouraging them to have children earlier, and it is possible that women in control states will eventually catch up, especially since very few women report a preference for having no children as their completed fertility.

had a boy. This is consistent with the policy's design, which only provides benefits for girls born after its implementation. Consequently, the increase in the 1B1G composition can now be more confidently attributed as the causal effect of the girl-boy policy.

Finally, as a falsification exercise, I analyze women who had exactly two children before the policy's implementation. A priori, we would not expect these women to be affected by the policy, as girls born before the policy are not retrospectively eligible for the benefits.[41] Additionally, any further children, whether boys or girls, would exceed the two-child limit, thereby disqualifying the parents from receiving benefits.[42] In Figure A2, I examine the effect of the policy on them having exactly two children i.e., discontinuing further childbearing. Women in this group display no differential behavior between treated and control states, reinforcing the validity of the policy's effects on groups not eligible for the benefits.

In summary, neither the girl-boy policy nor the girls-only policy increases the likelihood of women having only one daughter or two daughters, despite these compositions being eligible for the subsidies. This indicates that the current subsidy amount, though comparable to the cost differential between raising boys and girls, is insufficient to overcome the strong parental preference for at least one son, highlighting that this preference goes beyond mere economic considerations. Conversely, the girl-boy policy does encourage women to stop at a one-boy-one-girl composition, particularly among those who were previously childless or already had a son. These women seem content to forgo their desire for a larger family, satisfied with a smaller family that includes a son, who, despite not receiving a subsidy, fulfills their preference for at least one boy, and a daughter, who benefits from the policy's incentives.

Heterogeneity Based on Age at the Time of the Policy Implementation One crucial determinant of the fertility response is a woman's age, since fecundity declines over time (Noord-Zaadstra et al., 1991), and older women at the time of policy implementation were

[41] One exception is children born in 2006, who are eligible for benefits despite parents having two children at what I define as the start of the policy, potentially influencing stopping behavior.

[42] Another situation where parents could be affected is if their children passed away, returning them to a child composition that could make them eligible for benefits. For instance, consider a woman who had two daughters in the year the girl-boy policy was introduced, but one daughter died soon after. If she then gives birth to another girl, she may be eligible to register this newborn girl under the policy, as the newborn girl would have one living sibling at the time of registration.

less likely to be affected, having potentially already made fertility decisions that rendered them ineligible. For example, a 30-year-old woman in a treated state with three children at the time of policy implementation would exceed the two-child limit mandated under these policies and thus be ineligible for benefits. Thus, I estimate Equation (1.1) and Equation (1.2) separately for the following three age groups: (1) women who were aged 10-19 years at the end of τ, i.e., the year of the policy implementation, (2) women who were aged 20-29 years at the end of τ, and (3) who were aged 30-39 years at the end of τ.

Figures 1.9a and 1.9b display event study estimates by age group for the girl-boy and girls-only policies, respectively, with green representing women aged 10-19, blue representing those aged 20-29, and red for women aged 30-39. Table 1.7 provides the corresponding difference-in-differences (DID) estimates. The results indicate that the girls-only policy has insignificant effects[43] across all age groups. In contrast, the girl-boy policy leads to an increase in the proportion of women with targeted child compositions, with the effect being more pronounced among younger women aged 10–19[44] and 20–29 at the time of policy implementation. The probability rises by 1.78 percentage points for those aged 10–19 at the time of implementation and by 1.45pp for those aged 20–29. Among women aged 30–39, the estimate is smaller in magnitude (0.32pp) and statistically indistinguishable from zero.

I also verify that the estimates for both the policies, conditional on the age of the woman at the time of policy implementation, are of the same magnitude as in Columns 2-4 and 6-8 of Table A2 when using the border strategy, i.e., by restricting the analysis to bordering districts of the treated state. I also verify that the results look similar when using the synthetic control method, as evidenced in Figures A4a and A4b.

Heterogeneity Based on Wealth Since the policies target low-income households, I divide my sample into five wealth quintiles using a wealth index constructed from asset information collected during the surveys. The DID results (Figure 1.10a) show more pronounced effects in lower wealth quintiles of the girl-boy policy on the likelihood of women

[43]For women aged 30-39, the effect size is 1.01pp, but is statistically indistinguishable from zero as the standard errors are much wider. For the other two age groups, the effect sizes are smaller in magnitude as well as statistically insignificant.

[44]These women may appear very young to have children, but this refers to their age at the time of policy implementation. The analysis measures the impact of the policy up to seven years into the future, during which period these women age accordingly.

achieving a targeted child composition. On the other hand, the girls-only policy (Figure 1.10b) shows smaller and statistically insignificant impact across all wealth categories.

1.6.1.2 Sterilization at Targeted Compositions

Since the policies I study required sterilization at eligible child compositions for women to enroll their daughters in these programs, I now examine the effect of these policies on sterilization behavior among women at the targeted child compositions. Figure 1.11 presents the main results on the impact of the policies on sterilization at these compositions. The green line represents event study estimates for the girl-boy policy, while the blue line depicts estimates for the girls-only policy. Table 1.8 reports the corresponding DiD estimates. I find that both policies encouraged sterilization of women at targeted child compositions. The girl-boy policy increased this probability by 0.87pp (Column (1) of Table 1.8, relative to a mean sterilization rate of 3.6% in the control group, while the girls-only policy had an effect of 0.56pp (Column (2) of Table 1.8), relative to a control mean of 0.8%. These effects represent large relative increases of approximately 24% and 68%, respectively.

Sterilization, as a permanent and irreversible procedure,[45] carries potential disutility for women due to both the loss of future fertility options (a ”psychic cost” associated with the inability to reconsider their decision) and the risks inherent in the procedure, including complications and even death.[46] A possible explanation for the large observed increase in sterilization, even under the girls-only policy, which had a small (0.19pp) and statistically insignificant effect on the likelihood of women achieving a subsidy-eligible composition, is that women who would have otherwise stopped at 0B1G or 0B2G (due to a preference for smaller families[47] and no strong son preference) are now incentivized to undergo sterilization solely to secure the financial benefits for their daughters offered by the policy.

[45]While sterilization is theoretically reversible, it is seldom reversed in practice due to technical challenges and financial constraints.

[46]Short-term risks associated with the sterilization procedure include haemorrhage, anaesthesia related complications, trauma to blood vessels, infections, etc. These risks are compounded by widespread non-compliance with standard surgical protocols, particularly in sterilization camps and public hospitals across India (Bali et al., 2020; Koenig et al., 2000). Rupelle and Dumas (2025) finds that female sterilization increases the number of symptoms in gynecological health by 42% in India.

[47]The preference for smaller families could itself be motivated by health reasons, as higher order pregnanacies are more likely to be associated with maternal mortality (Rupelle and Dumas, 2025).

Heterogeneity based on Age at the Time of the Policy Implementation I examine the heterogeneity in the effects of the CCT policies on sterilization at subsidy-eligible child compositions based on women's age at the time of policy implementation. In Columns 1-3 of Table 1.9, I find that the sterilization effects under the girl-boy policy are concentrated among younger women (those under 30 at the time of implementation). The effect sizes for women aged 10-19 and 20-29 at the time of the policy implemenation are comparable: 1.09pp and 0.90pp, respectively. These effects are meaningful relative to baseline sterilization rates in the control states (0.64% and 5.38%, respectively), suggesting a substantial behavioral response. For women aged 30–39, the estimate is negative (-0.37pp) and not statistically significant. Figure 1.12a presents the event study version of this analysis. Under the girls-only policy, statistically significant increases are observed across all age groups (columns 3-6 of of Table 1.9). Women aged 10–19 experience a 0.88pp increase (relative to a baseline of 0.15%), those aged 20–29 a 0.25pp increase (relative to a 1.18% mean), and even the oldest age group shows a 0.59 pp increase (relative to a 2.20% baseline), significant at the 10% level. Figure 1.12b presents the event study version of this analysis.

Heterogeneity by Wealth In Figures 1.13a and 1.13b, I examine heterogeneity in the effects of the girl-boy and girls-only policies, respectively, on sterilization at subsidy-eligible child compositions based on women's wealth. I present the DiD coefficients and their confidence intervals in these figures for different subsamples based on wealth quintiles. As expected, I find stronger effects for less wealthy women.

Sterilization at Non-Targeted Compositions To confirm that the increase in sterilization at targeted compositions is driven by the subsidies and not by broader changes in sterilization practices across states or over time, I examine sterilization rates at non-targeted compositions (Figure A5). The analysis shows that there are no discernible pre-trends, and the post-policy coefficients are both statistically insignificant and close to zero. This supports the conclusion that the observed increase in sterilization is directly attributable to the policies.

Sterilization at Targeted Compositions within Policy Prescribed Time The policies require sterilization to occur within specific time limits in order for the family to qualify for the benefits: within one year for the girl-boy policy and within three years for the girls-only policy. The overall (Figure A6, Table A3) and the age-group specific results (Table A4, Figure A7a for the girl-boy policy, and Figure A7b for the girls-only policy) remain qualitatively similar when analyzing sterilizations that occurred within these policy-prescribed time limits, although the effect sizes are somewhat smaller. This finding is not particularly surprising, as most sterilizations tend to happen within 1 to 3 years of childbirth.

1.6.1.3 Number of Children

For the girl-boy policy, Figures 3.4b and 3.4c showed a reduction in the proportion of women who go on to have more than two children, as well as a decline in the percentage of women who remain childless. To capture the net effect on fertility, I also analyze the total number of children a woman has by the end of year t. The results (Figure 1.14 and Table 1.10) indicate that the girl-boy policy leads to a decrease in the total number of children by 3.6% (a decrease of 0.0419 children relative to the baseline mean of 1.16 children per woman), while the girls-only policy leads to a modest decrease of 0.0052 children, which is statistically insignificant at the 10% level. The small impact on total fertility under the girls-only policy might seem surprising, particularly since sterilization is one of most foolproof methods to prevent unintended pregnancies. Byker and Gutierrez (2021) found that access to sterilization in Peru led to an average decrease in completed fertility of approximately 0.8 children per woman. Assuming a similar effect size for sterilization in this context, the estimated reduction in unintended pregnancies in my study falls within the confidence intervals of the observed fertility effect. However, the reduction in the number of children observed under the girl-boy policy exceeds what would be expected from sterilization preventing unintended pregnancies alone. This can be explained by the fact that women are actively choosing to limit their family size under this policy, as evidenced by the increase in the 1B1G composition and the decline in higher parity births.

1.6.2 Effect on Post-Birth Outcomes for Girls

As mentioned in Section 1.2, there is often a gender disparity in post-birth outcomes, with some households neglecting and substantially under-investing in their daughters' health and education. For instance, infant and under-5 child mortality rates are higher for girls, as parents tend to neglect their daughters, particularly during infancy and early childhood, a crucial period for development. The staggered payment structure under the policies I study is designed to combat this issue, as girls need to survive and fulfill certain education milestones to receive the financial benefits. To analyze the impact of these policies, I turn to the child-outcome dataset that I construct for policy-eligible girls. As a reminder, I define policy eligibility at birth: girls born in post-program years are eligible if they have at most one sibling in the girl-boy policy, and at most one sister (but no brothers) at birth under the girls-only policy.

1.6.2.1 Effect on Mortality

Figures A8a and A8b, along with the corresponding Difference-in-Differences (DiD) table (Table A5), show no evidence of differential trends between treated and control states prior to and after the implementation of either policy, which indicates that, despite the staggered payments tied to child survival in both policies, there was no measurable impact on reducing infant or under-5 mortality rates for policy-eligible girls.

1.6.2.2 Effect on Enrollment in School

The girls-only policy includes an important enrollment-based conditionality, requiring that girls be admitted to school by age five to retain eligibility for benefits. My results (Figure 3.3 and Table 3.3) show that the girls-only policy successfully promotes earlier school enrollment by mandating that girls start school one year before the legally mandated age of six. There is a 3.86pp, or a 5.9% increase relative to the baseline of 65% children being enrolled. Thus, even though the girls-only policy does not significantly shift fertility outcomes and inefficiently incentivizes sterilization, it nonetheless illustrates how incorporating education-based conditionalities into primarily fertility-focused programs can influence behavior. Investments in early childhood education are known to be critical for improving long-term outcomes (Cunha and Heckman, 2007). However, I am unable to examine such long-term effects, such as high school completion or adult income. The available data only extend through 2019, which would be insufficent to look these outcomes, even for the first cohort of girls exposed to the policy (beginning in 2006). Further research will be needed to assess these longer-run impacts.

1.7 Framework

As previously discussed, son preference in India arises from a variety of economic and sociocultural factors, including dowry practices that increase the perceived cost of raising daughters, religious norms, patrilineality, and patrilocality. In Section 2.3, I showed that son preference influences fertility behavior, with women making sequential childbearing and sex selection decisions based on factors such as income and the sex composition of their existing children. It is also clear from the data that the pattern of fertility behavior is consistent with a premium on having at least one son, which differs from a general daughter aversion. However, obtaining a son does not lead most parents to cease childbearing altogether; it typically reduces their likelihood of continuing but does not eliminate it, as many couples still hold a preference for larger family sizes. Recent survey evidence on fertility preferences in Haryana, a state with pronounced patriarchal institutions and one of the most skewed sex ratios at birth, confirms this: parents generally desire at least one son but, upon securing that, often prefer a balanced sex composition (Jayachandran, 2017).

Thus, if one were to write down a formal framework of fertility behavior that could explain the data patterns observed for poorer women, the framework should incorporate the following factors: the premium placed on having at least one son, preferences for family size, preferences for sterilization, and consumption utility. Not all women would have the same preferences though; for example, Norling (2018) demonstrates that there are combinations of preferences over the sex and number of children that explain observed childbearing, rather than assuming homogeniety in preferences within the population. Such a framework could help us understand who is on the margin of being affected by these policies, under different combinations of parameter values. This would also help rationalize the findings from the empirical section, which are threefold. First, neither policy affects the probability of having girls-only compositions. Second, there is an increase in the probability of a woman having a one-boy-one-girl composition, rather than continuing to have more than two children. Third, sterilization at subsidy-eligible child compositions increase under both policies.

Suppose the utility function of women is as follows:

$$u = \nu_1 1\{n = 1\} + \nu_2 1\{n = 2\} + \nu_3 1\{n \geq 3\} + \nu_4 1\{n^b \geq 1\} + \nu_5 1\{s = 1\} + (I - n^b c - n^g cx)$$

Here, n denotes the total number of children, with $n = n^b + n^g$, i.e., the sum of boys and girls. The term s is an indicator for sterilization, I denotes income, and c is the cost of raising a boy. The cost of raising a girl is given as $c^g = c \cdot x$, where $x > 1$ implies that girls are costlier to raise than boys. Since the policies under consideration disqualify families with more than two children, the main interesting margin is the transition from two to three children, and thus, I do not consider preference for more than three children. The parameter ν_4 represents a utility shifter from having at least one son. The parameter ν_5 captures the disutility associated with sterilization, which can represent both the physical and psychological costs of opting for a permanent method of family size control. Lastly, women receive a utility from consumption, with total consumption given by the residual income after deducting the costs of raising boys and girls.

Given a particular sex composition of their existing children, women make two sequential decisions: first, whether to continue childbearing, and second, whether to undergo sterilization. Sterilization acts as a binding decision, as it prevents any future childbearing. I assume risk neutrality for parents and assume they can perfectly control the number of children they have, even without sterilization. Additionally, I assume that there is no access to sex-selective abortions; thus, the sex of each subsequent child is essentially random (0.5 probability), a reasonable assumption for poorer populations where sex-selective abortions are less common. I further assume that the parameters ν_1, ν_2, ν_3, ν_4, are all non-negative, the parameter ν_5 is negative, and that all women have the same income and costs of raising children. Lastly, I assume that $\nu_1 > 0.5c(1+x)$ and $\nu_2 - \nu_1 > 0.5c(1+x)$, i.e., in the absence of subsidies, it would be optimal to have at least two children without sterilization.

One limitation of this framework is that no woman would choose sterilization without subsidy payments, as sterilization carries only disutility in this model. In reality, women may derive a net positive utility from an almost foolproof method of preventing future pregnancies, viewing it as a means of achieving their desired family size with greater certainty. However, this omission may not be entirely inconsistent with the context, as India offers one-time financial incentives to encourage sterilization among women.

Case I: Women with preference for larger families and with a strong son preference Suppose $\nu_3 - \nu_2 > 0.5c(1+x)$ and $\nu_4 > 2\{(\nu_3 - \nu_2) - \nu_5 + 0.5c(1+x)\}$

In the absence of any subsidies, the optimal decision that a woman faces at any existing composition will be based on the difference in the utility from stopping at that composition and continuing to have another child. Suppose the woman has two children. Then, the optimal decision will be to continue to have a third child and not sterilize, because $\nu_3 - \nu_2 > 0.5c(1+x)$. Now suppose there is a policy that gives a subsidy S to women who have a 0B1G, 0B2G, or 1B1G compositions as their completed fertility, conditional on sterilization. I will treat the subsidy amount as an income transfer to the parents for simplicity.

Woman is at 1B1G The amount of incentive, $S > \nu_3 - \nu_2 + 0.5c(1+x) - \nu_5$, i.e., the sum of disutility from stopping at a smaller family and sterilizing.

Woman is at 0B2G, with the woman being potentially eligible for a subsidy of 2S as she has two girls The amount of incentive, $S > 0.5\{\nu_3 - \nu_2 + 0.5c(1+x) - \nu_5 + 0.5\nu_4\}$, i.e., the sum of disutility from stopping at a smaller family, sterilizing, and not having a boy as a potential third child. Because of the assumption that son preference is strong, the amount of subsidy required to shift women's fertility behavior when they already have a son is lower than the amount of incentive required to switch to an all-girls composition.

Case II: Women with preference for smaller families and with no son preference Suppose $\nu_3 - \nu_2 < 0.5c(1+x)$ and $\nu_4 = 0$

In the absence of subsidies, women who have two kids would stop without sterilization since the third child will decrease their utility. Subsidies of size $S > -\nu_5$ are required for the woman to now sterilize at two children, irrespective of whether that composition is 1B1G or 0B2G.

Thus, this framework suggests that the marginal woman affected by the policy depends crucially on her underlying preferences for family size and at least one son, and her tolerance or aversion to sterilization. This heterogeneity in these three parameter values helps reconcile the empirical result that both the girl-boy and girls-only policies increased sterilization at subsidy-eligible compositions, but only the former influenced child composition. For women with strong son preference and larger family desires, only relatively modest subsidies are needed to encourage sterilization once they have at least one son secured (e.g., 1B1G). In contrast, the same subsidy may not suffice to adopt sterilization and deter further child-bearing among women without a son. An effective fertility-targeted CCT should, therefore, account for the underlying fertility ideals, son preference, and sterilization disutility.

1.8 Conclusion

In this paper, I evaluate two large CCT programs that provide generous subsidies for girls in smaller families, one that is restricted to families without boys and one without such a restriction. Both policies require parental sterilization, and payments are staggered. I find that the policy without the restriction to girls increases the probability of having a subsidy-eligible child composition by 8%, driven by an increase in women with a one-boy-one-girl composition. This effect is concentrated among younger women (aged 10-29 at the time of the policy implementation), poorer women, women with no prior children, and those with a prior boy. Overall fertility decreases, with the number of children decreasing by 3.4%. In contrast, all-girl compositions do not increase under either policy, reflecting a premium for the first son, that can't be compensated by the current subsidies. Both policies increase sterilization at subsidized compositions, but the girl-restricted policy does so without affecting fertility, suggesting that it inefficiently incentivizes women to undergo the risky procedure. However, the girl-restricted policy, by explicitly requiring subsidized girls to start school one year before the legally mandated age, increased early school enrollment, highlighting the importance of education-based conditionalities in these transfer programs.

My findings contribute to a broader discussion on designing an optimal policy aimed at addressing both the birth and well-being of girls in smaller families. In contexts where son preference is driven not only by the cost differential between raising girls versus boys but also by deep-rooted social norms, policies targeting fertility decisions and outcomes of girls must be carefully calibrated. My results suggest that the demand elasticity for girls is higher in the presence of a son, and thus, a key consideration in designing such policies is how flexible they are in terms of the number of boys allowed, particularly if increasing subsidy amounts is not feasible from a public expenditure perspective.

1.9 Figures and Tables

Figures

Figure 1.1: Marital and Fertility Outcomes by Age

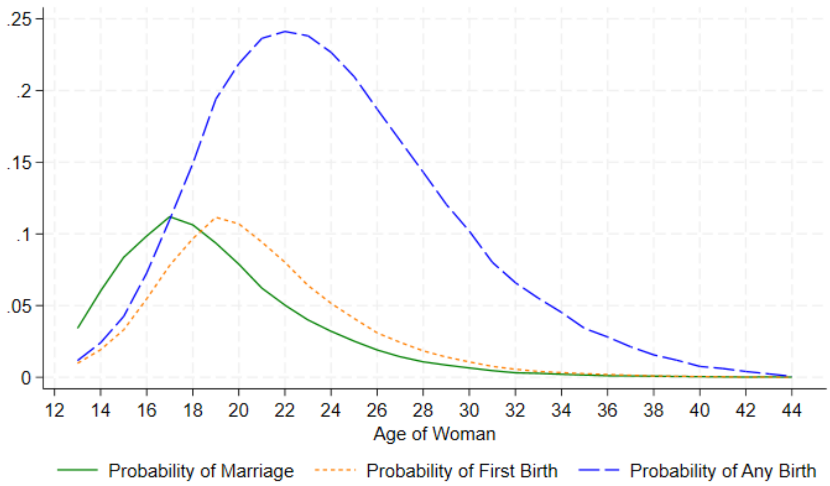

(a) Probability of Marriage, First Birth, and Any Birth by Age

Note: This figure plots the probability of marriage, first birth, and any birth, at different ages for women in India. Data source: DHS.

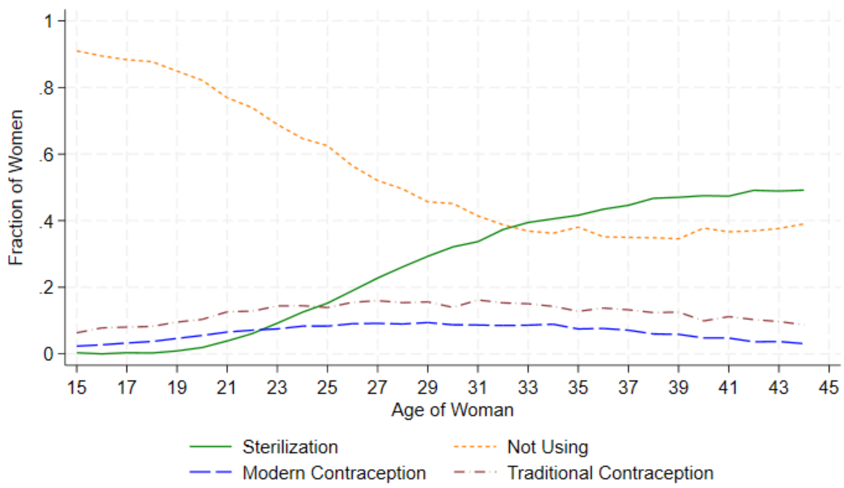

(b) Contraception Use by Age

Note: This figure displays the percentage of women, stratified by age, reporting the use of various contraceptive methods at the time of the survey. Traditional methods include rhythm/periodic abstinence, lactational amenorrhea, withdrawal, etc. Data source: DHS.

Figure 1.2: Trends in Total Fertility Rate and Child Sex Ratio (0-6 years) in India

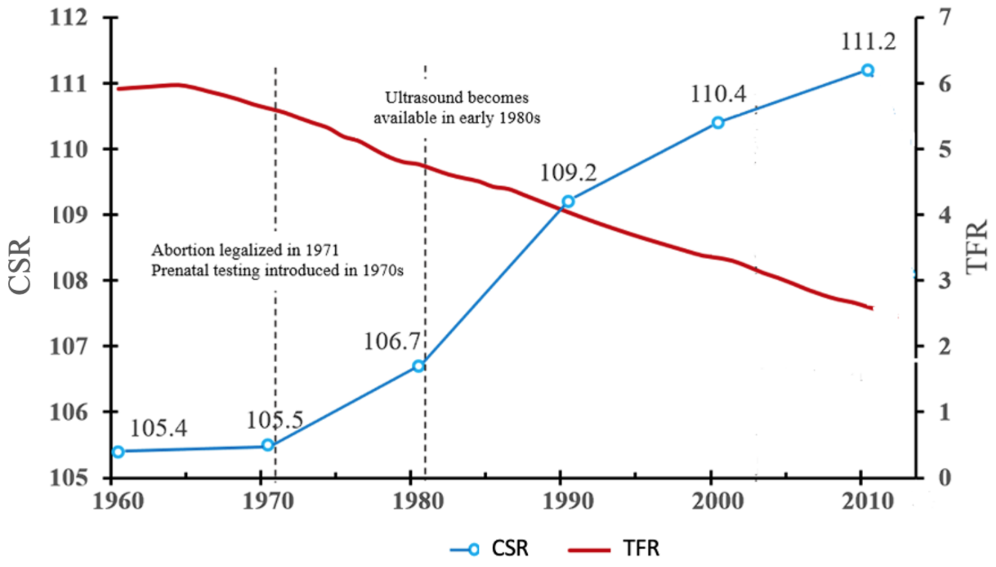

Note: This figure plots the trends for Total Fertility Rate (TFR) and Child Sex Ratio (CSR) for children aged 0-6 years. Data for TFR come from UNDP and data for CSR are sourced from Census of India.

Figure 1.3: Sex Composition of Children of Women With Completed Fertility

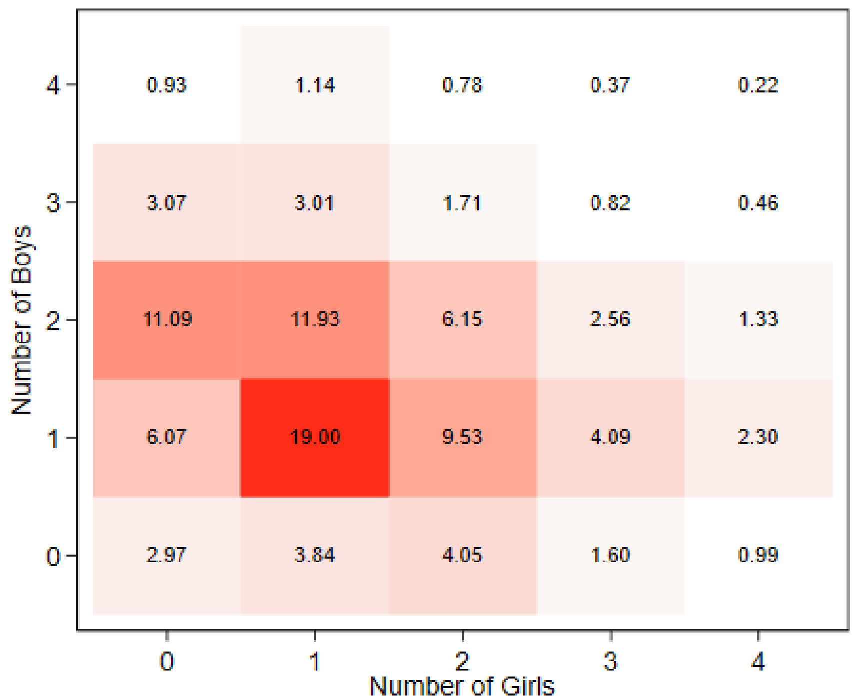

Note: This figure shows the percentage of women aged 30+ in 2015-16 who have a particular sex composition of children.

Figure 1.4: Parity Progression Ratio

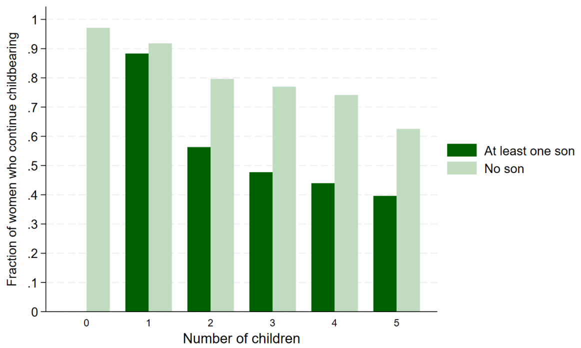

(a) Overall

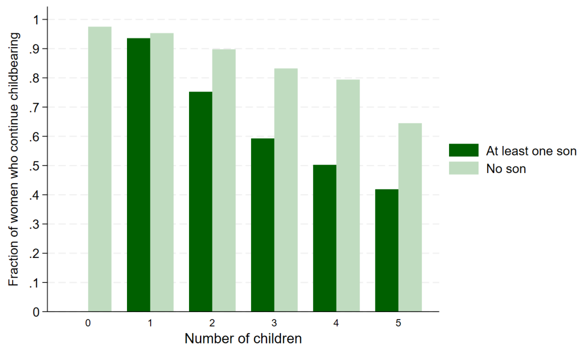

(b) Least Wealthy

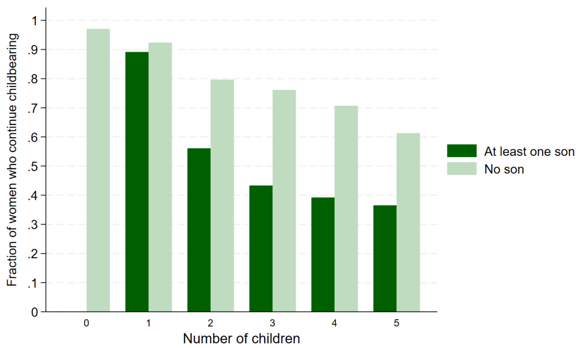

(c) Moderately Wealthy

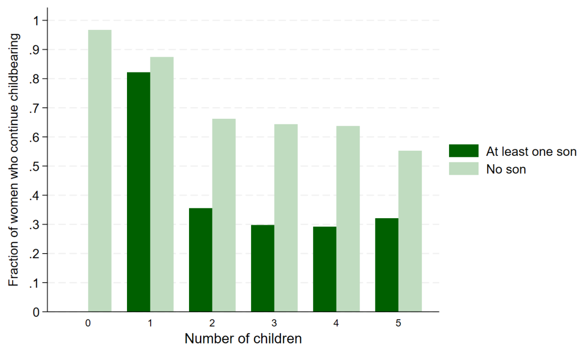

(d) Most Wealthy

Note: This figure shows the probability of a woman continuing childbearing, conditional on the sex composition of her existing children and her wealth class. Data is for women aged 30+ in 2015–16, sourced from the DHS.

Figure 1.5: Probability Sex of Next Child is Female by Wealth

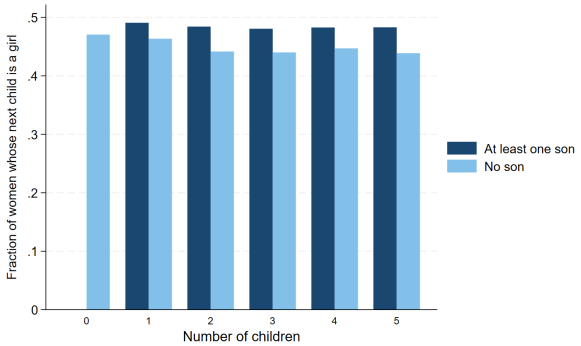

(a) Overall

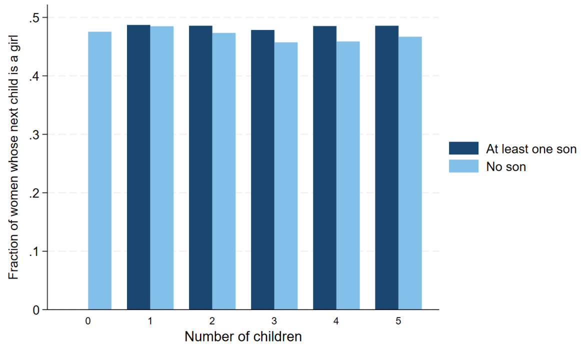

(b) Least Wealthy

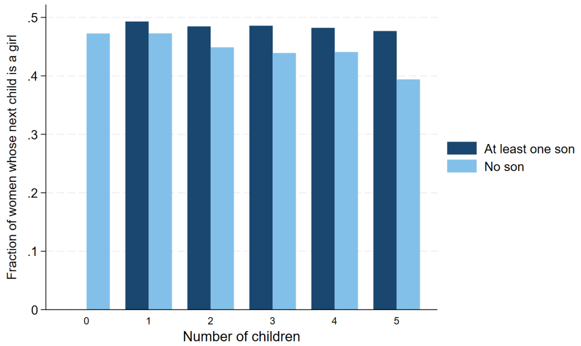

(c) Moderately Wealthy

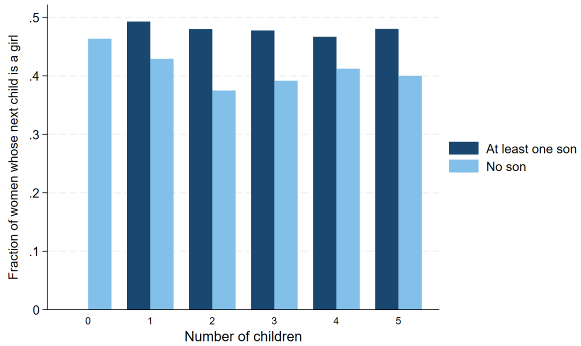

(d) Most Wealthy

Note: This figure shows the probability of a woman giving birth to a girl as her next child, conditional on the sex composition of her existing children and her wealth class. Data is for women aged 30+ in 2015-16, sourced from the DHS.

Figure 1.6: Treatment and Control States

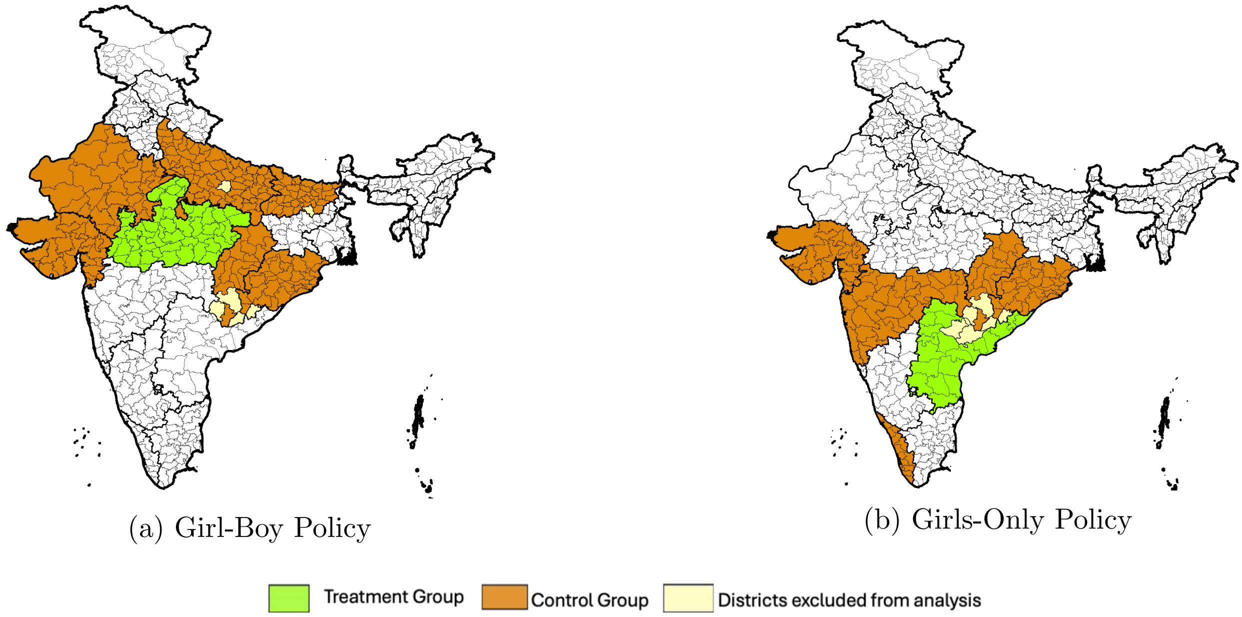

Note: This figure shows the treatment and control states for the girl-boy and the girls-only policy. The treatment state is indicated by green and the control states by orange. States in white are not considered as control states and the districts in yellow are excluded from the treatment or control states due to the presence of a different girl-promotion policy during the time period being studied.

Figure 1.7: Effect of Policies on Subsidy-Eligible Compositions

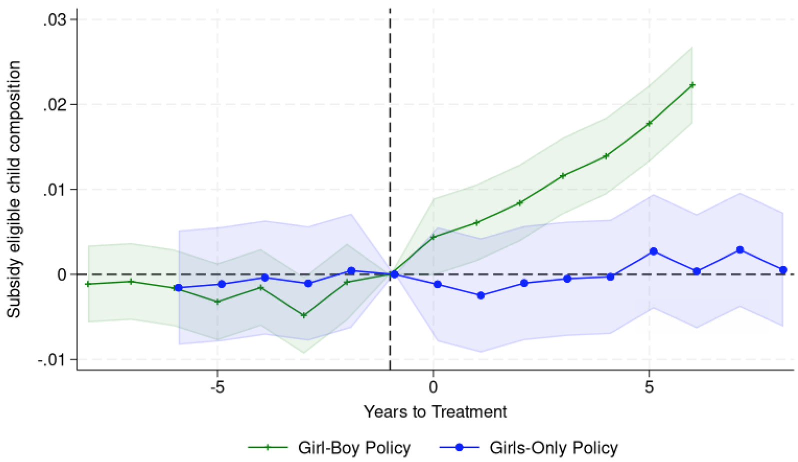

Note: This figure shows the effect of the Girl-Boy (green) and girls-only (blue) policies on the probability that a woman has a subsidy-eligible child composition. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.8: Effects of Girl-Boy Policy on Different Child Compositions

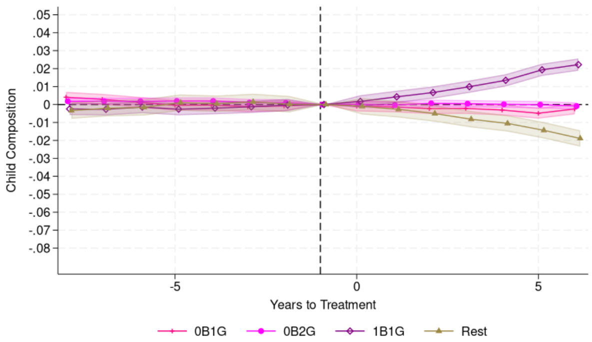

(a) Overall sample

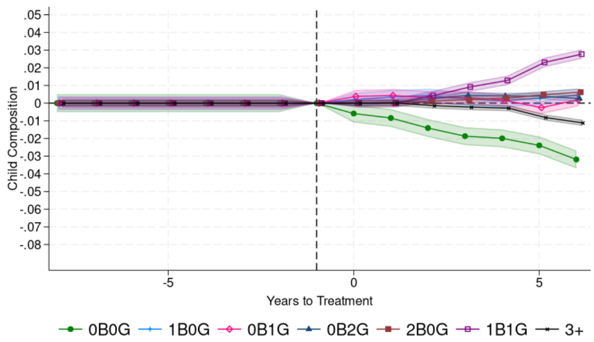

(b) Women with zero kids pre-policy

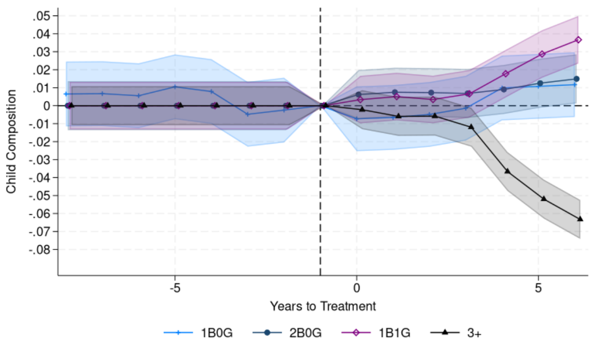

(c) One Boy

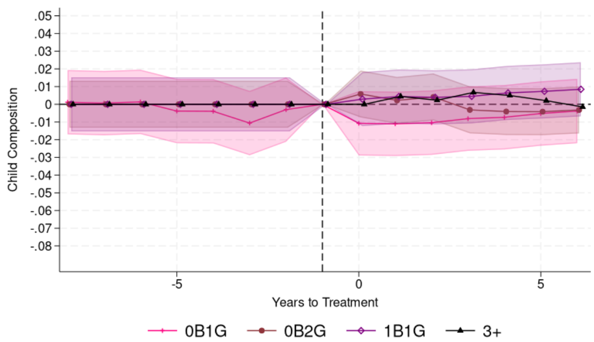

(d) One Girl

Note: This figure shows the effect of the Girl–Boy policy on the probability that a woman has different child compositions. Subfigure (a) presents results for the full sample, while subfigures (b)–(d) restrict the sample to women who, in the first year of policy implementation, had zero children, one boy, and one girl, respectively. Each panel plots the estimated β_z coefficients from the event study specification in equation (1.2), along with their 95% confidence intervals. Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.9: Effect of Policies on Subsidy-Eligible Child Compositions By Age of Woman at Time of Policy Implementation

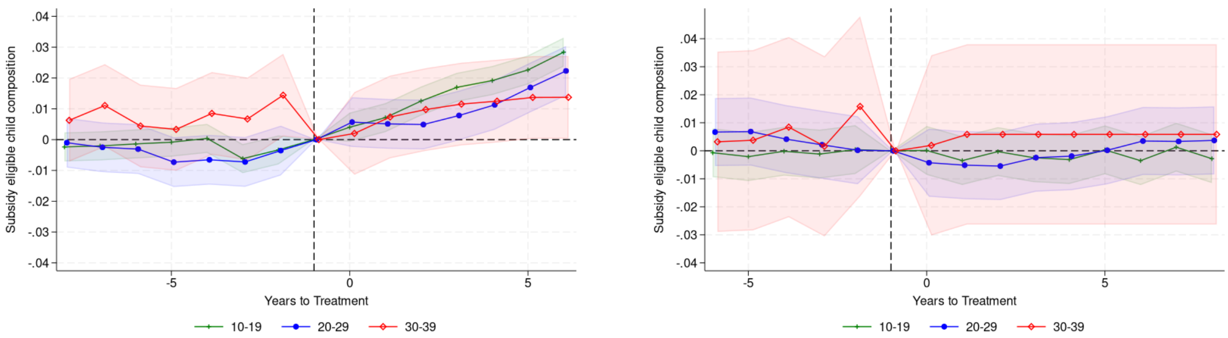

(a) Girl-Boy Policy (b) Girls-Only Policy

Note: This figure shows the effect of the Girl-Boy and the Girls-Only policy on the probability that a woman has a subsidy-eligible child composition, stratified by age group in the first year of policy implementation. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, and urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.10: Effect of Policy on Subsidy-Eligible Child Compositions By Wealth

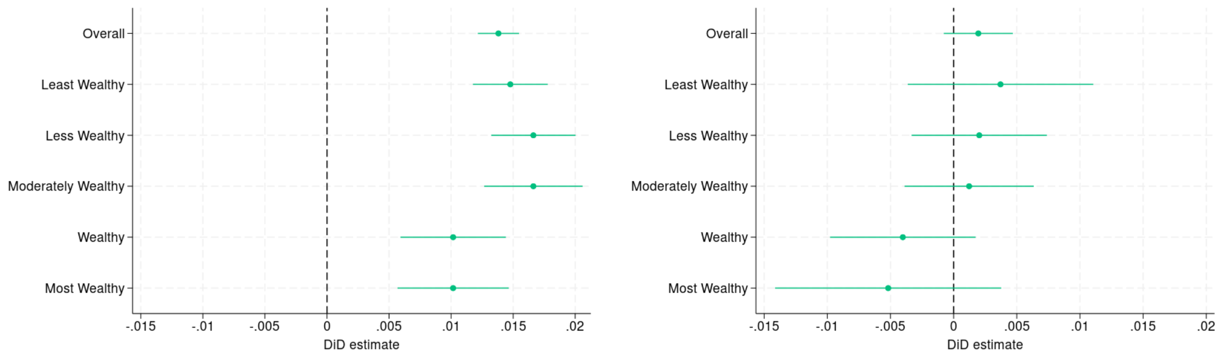

(a) Girl-Boy Policy (b) girls-only Policy

Note: This figure shows the effect of the Girl-Boy (panel (a)) and girls-only (panel (b)) policies on subsidy-eligible child compositions, disaggregated by wealth levels. Each panel plots the β coefficient and its 95% confidence interval derived from estimating the difference-in-differences (DiD) equation (1.1). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap. A wealth index is constructed using principal component analysis on variables such as a household's ownership of selected assets, such as televisions and bicycles; materials used for housing construction; and types of water access and sanitation facilities. All these variables are measured at the time of the survey. Individuals are categorized into five wealth quintiles based on their relative standing within each survey round.

Figure 1.11: Effect of Policies on Sterilization at Subsidy-Eligible Compositions

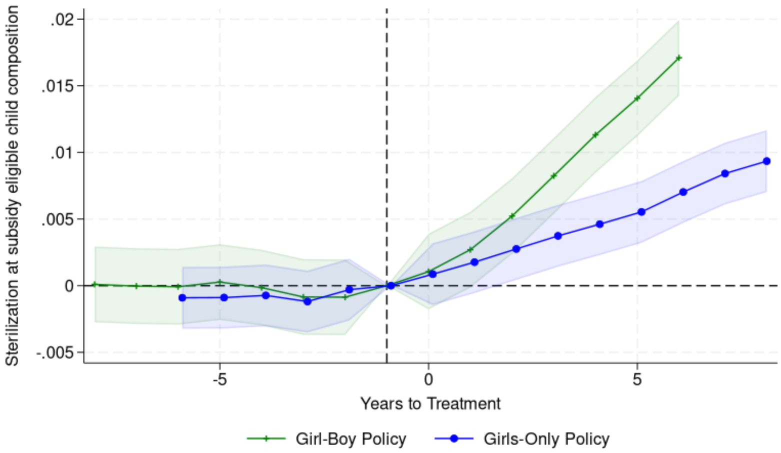

Note: This figure shows the effect of the Girl-Boy (green) and girls-only (blue) policies on the probability of sterilization at a subsidy-eligible child composition. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.12: Effect of Policies on Sterilization at Subsidy-Eligible Child Compositions By Age of Woman at Time of Policy Implementation

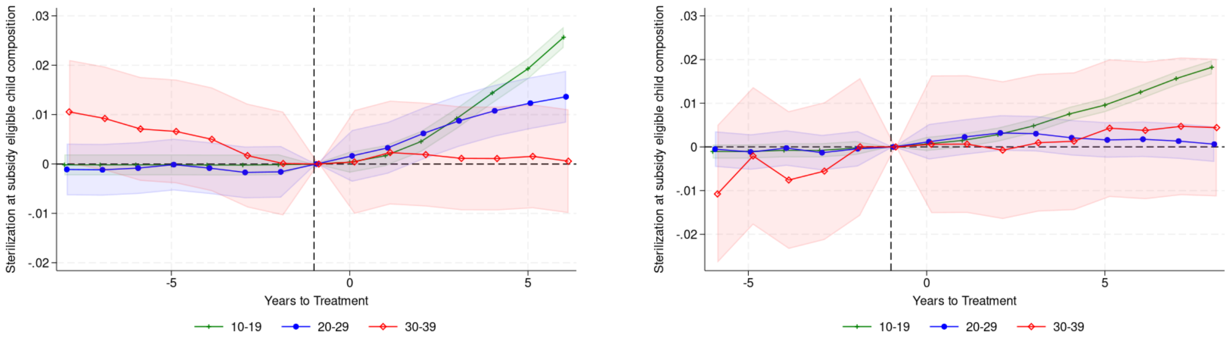

(a) Girl-Boy Policy (b) Girls-Only Policy

Note: This figure shows the effect of the Girl-Boy and the Girls-Only policy on the probability that a woman is sterilized at a subsidy-eligible child composition, stratified by age group in the first year of policy implementation. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, and urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.13: Effect of Policy on Sterilization at Subsidy-Eligible Child Compositions By Wealth

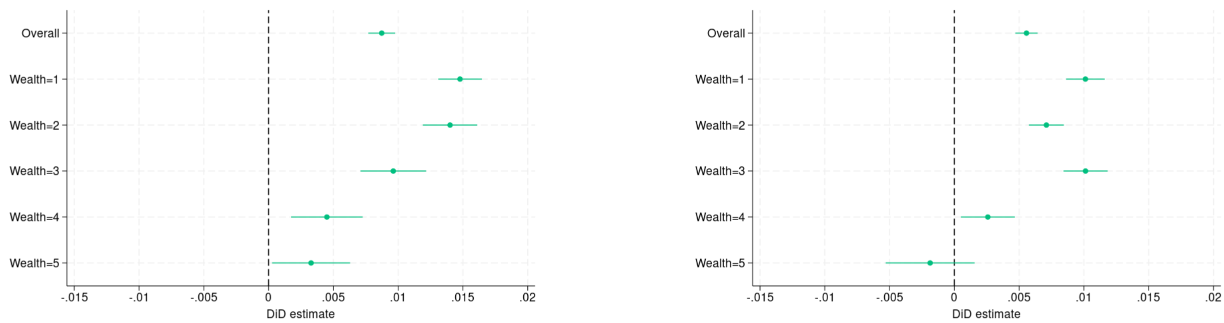

(a) Girl-Boy Policy (b) Girls-only Policy

Note: This figure shows the effect of the Girl-Boy (panel (a)) and Girls-only (panel (b)) policies on sterilization at subsidy-eligible child compositions, disaggregated by wealth levels. Each panel plots the β coefficient and its 95% confidence interval derived from estimating the difference-in-differences (DiD) equation (1.1). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap. A wealth index is constructed using principal component analysis on variables such as a household's ownership of selected assets, such as televisions and bicycles; materials used for housing construction; and types of water access and sanitation facilities. All these variables are measured at the time of the survey. Individuals are categorized into five wealth quintiles based on their relative standing within each survey round.

Figure 1.14: Effect of Policies on Number of Children

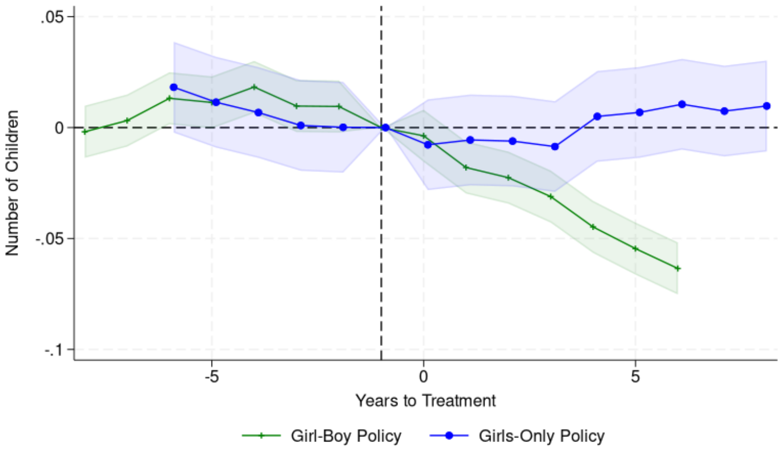

Note: This figure shows the effect of the Girl-Boy (green) and girls-only (blue) policies on the number of children. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation 1.2. Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure 1.15: Effect of Girls-Only Policy on Enrollment in School by Age-5 of Policy-Eligible Girls

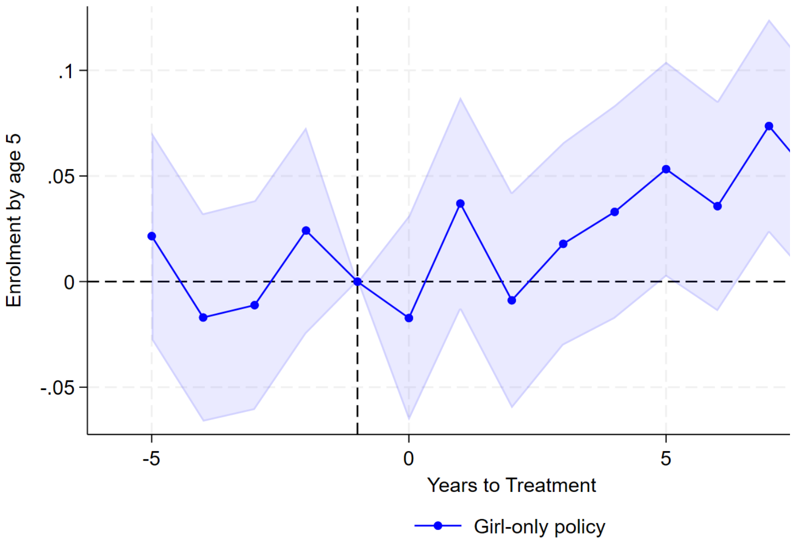

Note: This figure shows the effect of the Girls-Only policy on the probability of policy-eligible girls to be enrolled in school by age five. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states and birth years. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Tables

Table 1.1: Payment Structure under Girl-Boy Policy

Condition	**Benefit**
Admission in Grade 6	Rs. 2000 ($50)
Admission in Grade 9	Rs. 4000 ($100)
Admission in Grade 11	Rs. 7500 ($187.50)
During Grade 11 and 12	Rs. 200 per month ($5 per month)
Final Payment Condition	**Terminal Benefit**
At age 21, conditional on:	
• Appearance in Class 12th Exams	Rs. 100,000 ($2500)
• Staying unmarried till age 18	
Present Discounted Value of Total Benefits	Rs. 65,300 ($1632.50)

Note: Present Discounted Value (PDV) is calculated assuming a discount rate of 3%. An exchange rate of 40Rs/$ is assumed. It is not necessary for the girl to graduate; she just has to appear in the Grade 12 examinations to be eligible to receive the terminal payment.

Table 1.2: Payment Structure under Girls-Only Policy

Eligibility conditions after enrollment	
• Enrolled in school by age 5	
Condition	**Benefit**
Admission in Grade 9 -12	Rs. 1200 ($30) per year
Final Payment Condition	**Terminal Benefit**
At age 20, conditional on:	
• Appearance in Class 12th Exams	Rs. 100,000 ($2500) if one girl
• Staying unmarried till age 18	Rs. 30,000 ($750) if two girls
Present Discounted Value of Total Benefits	
One Girl family	Rs. 58,405 ($1460)
Two Girl family	Rs. 19,648 ($491)

Note: Present Discounted Value (PDV) is calculated assuming a discount rate of 3%. An exchange rate of 40Rs/$ is assumed.

Table 1.3: Summary Statistics for Girl-Boy Policy

Variable	**Control States**		**Treated State**	
	Mean	**SD**	**Mean**	**SD**
Eligible child composition	0.172	(0.340)	0.157	(0.347)
Sterilized at eligible child composition	0.036	(0.100)	0.028	(0.127)
Urban resident	0.214	(0.410)	0.252	(0.434)
Education (in years)	5.520	(5.280)	5.338	(4.933)
Hindu	0.881	(0.324)	0.926	(0.262)
Other Backward Caste	0.495	(0.500)	0.452	(0.498)
Scheduled Caste	0.196	(0.397)	0.159	(0.366)
Scheduled Tribe	0.128	(0.334)	0.223	(0.416)
Wealth Index (scale: 1-5)	2.650	(1.396)	2.612	(1.426)
Age at start of policy (in years)	21.395	(6.641)	21.526	(6.602)
Observations	2,258,088		478,368	

Note: This table shows the summary statistics of different variables in the control states and treatment state for the Girl-boy policy. Data are restricted to pre-treatment years.

Table 1.4: Summary Statistics for Girls-Only Policy

Variable	**Control States**		**Treated State**	
	Mean	**SD**	**Mean**	**SD**
Eligible child composition	0.107	(0.272)	0.083	(0.275)
Sterilized at eligible child composition	0.008	(0.059)	0.012	(0.110)
Urban resident	0.283	(0.450)	0.253	(0.435)
Education (in years)	7.052	(4.967)	6.511	(5.078)
Hindu	0.867	(0.340)	0.458	(0.498)
Other Backward Caste	0.399	(0.490)	0.227	(0.419)
Scheduled Caste	0.143	(0.350)	0.112	(0.315)
Scheduled Tribe	0.223	(0.416)	0.518	(0.500)
Wealth Index (scale: 1-5)	2.978	(1.400)	2.899	(1.182)
Age at start of policy (in years)	20.657	(6.337)	20.015	(6.011)
Observations	804,438		89,436	

Note: This table shows the summary statistics of different variables in the control states and treatment state for the Girls-Only policy. Data are restricted to pre-treatment years.

Table 1.5: Effect of Policies on Subsidy-Eligible Compositions (DiD estimates)

	Girl-Boy	Girls-Only
Treat × Post	0.0138***	0.0019
	(0.0008)	(0.0014)
Mean (Control States)	0.1721	0.1074
Observations	5,130,855	2,234,685

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy and girls-only policies on the probability of a woman having a subsidy-eligible child composition. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.6: Effect of Girl-Boy Policy on Different Child Compositions (DiD estimates)

	One Girl	One Boy One Girl	Two Girls	Rest
Treat x Post	−0.0010	0.0155***	−0.0006	−0.0137***
	(0.0006)	(0.0006)	(0.0004)	(0.0008)
Mean (Control States)	0.0623	0.0835	0.0264	0.8279
Observations	5,130,855	5,130,855	5,130,855	5,130,855

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy policy on the probability of a woman having different child compositions. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.7: Effect of Policies on Subsidy-Eligible Child Compositions by Age at Time of Policy Implementation (DiD estimates)

	Girl-Boy Policy			Girls-Only Policy		
	10–19	20–29	30–39	10–19	20–29	30–39
Treat × Post	0.0178***	0.0145***	0.0032	−0.0022	−0.0042	0.0101
	(0.0009)	(0.0015)	(0.0025)	(0.0017)	(0.0026)	(0.0062)
Mean (Control States)	0.0776	0.2481	0.2292	0.0776	0.1420	0.1075
Observations	2,200,425	2,208,960	721,470	1,017,255	1,001,460	215,970

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy and Girls-Only policies on the probability that a woman has a subsidy-eligible child composition, stratified by her age group at the time of policy implementation. Each coefficient is estimated using equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.8: Effect of Policies on Sterilization at Subsidy-Eligible Compositions (DiD estimates)

	Girl-Boy	Girls-Only
Treat × Post	0.0087***	0.0056***
	(0.0005)	(0.0004)
Mean (Control States)	0.0362	0.0082
Observations	5,130,855	2,234,685

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy and girls-only policies on the probability of a woman sterilizing at a subsidy-eligible child composition. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.9: Effect of Policies on Sterilization at Subsidy-Eligible Compositions by Age at Time of Policy Implementation (DiD Estimates)

	Girl-Boy Policy			Girls-Only Policy		
	10–19	20–29	30–39	10–19	20–29	30–39
Treat × Post	0.0109***	0.0090***	−0.0037	0.0088***	0.0025***	0.0059*
	(0.0004)	(0.0010)	(0.0023)	(0.0003)	(0.0008)	(0.0031)
Mean (Control States)	0.0064	0.0538	0.0794	0.0015	0.0118	0.0220
Observations	2,200,425	2,208,960	721,470	1,017,255	1,001,460	215,970

Note: This table shows the point estimate and the standard error of the effect of the Girls-Boy and the Girls-Only policy on the probability of a woman sterilizing subsidy-eligible child composition, stratified by the age group of the woman in the first year of policy implementation. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.10: Effect of Policies on Number of Children (DiD Estimates)

	Girl-Boy	Girls-Only
Treat × Post	−0.0419***	−0.0052
	(0.0021)	(0.0039)
Mean (Control States)	1.1639	1.0834
Observations	5,130,855	2,234,685

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy and girls-only policies on the number of children. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table 1.11: Effect of Girls-Only Policy on Enrollment in School by Age 5 (DiD Estimates)

	Girls-Only
Treat × Post	0.0386***
	(0.010)
Mean (Control States)	0.65
Observations	43,321

Note: This table shows the effect of the Girls-Only policy on the probability of enrollment in school by age 5 of policy-eligible girls. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states and birth years. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Chapter 2

Hiring Subsidies for the Disadvantaged: Evidence from the Work Opportunity Tax

2.1 Abstract

The US spends billions of dollars annually on policies to promote upward economic mobility. This paper studies the effectiveness of one such approach: targeted wage subsidies. We examine a large federal program, the Work Opportunity Tax Credit (WOTC), which covers up to 40% of wages for certain categories of disadvantaged workers and subsidizes over two million hires annually. Using linked administrative data, we apply four complementary quasi-experimental approaches to assess WOTC's impact on labor market outcomes, including age-based eligibility discontinuities, eligibility expansions, and changes to program participation costs. We find precise null effects on hiring, employment, and earnings across all specifications, where we can rule out effects on hiring as small as 0.2 percentage points. Instead, employers hire the same workers as in the absence of subsidies, and subsidies operate as a pure transfer to the firms, where most of the benefits accrue to a small number of large firms. To understand these results, we collect original data on the hiring processes of WOTC-utilizing firms. These data suggest that firms often silo information on candidates' subsidy eligibility from managers making hiring decisions, with some evidence that this is due to concerns about legal liability.

2.2 Introduction

The US spends billions of dollars each year on policies to promote employment and upward economic mobility among disadvantaged groups. Much of these efforts go to programs and policies focused on providing incentives to individuals to find work, such as the Earned Income Tax Credit and vocational training programs (Barnow and Smith, 2015). An alternative demand-side approach is wage subsidies, in which the government subsidizes the wages paid to certain types of workers to encourage firms to hire them (Katz, 1996).

This paper studies the largest such program in the US, the federal Work Opportunity Tax Credit (WOTC), which provides wage subsidies to encourage the hiring of groups typically disadvantaged in the labor market. Firms can apply for WOTC when making new hires and receive tax credits for up to 40% of the hire's wages in their first year in a job. WOTC was created in 1996 and initially subsidized the wages of approximately 400,000 individuals annually, most of whom were recipients of cash welfare. Over time, eligibility has expanded and now includes ten categories of individuals, the largest of which is Supplemental Nutrition Assistance Program (SNAP) recipients. The number of jobs subsidized by WOTC has grown to over 2.5 million annually (Department of Labor, 2023b), and in 2022, approximately 4% of all new hires in the US had their wages subsidized by WOTC (Bureau of Labor Statistics, 2023b). That is more than double the number of workers earning at or below the federal minimum wage (Bureau of Labor Statistics, 2023a), a policy that has received considerably more academic and policy attention.

In standard economic models, wage subsidies increase employment or wages by lowering the effective cost of labor to employers. If targeted to particular groups, they may encourage the hiring of workers from those groups by lowering the cost of these workers as compared to workers ineligible for subsidies. Temporary wage subsidies like WOTC can encourage firms to accept the risks of hiring workers whom they might otherwise be unwilling to hire (Pallais, 2014). Targeted wage subsidies may also affect non-employment outcomes, such as reducing reliance on public benefit programs or the likelihood of engaging in criminal activity.

While promising in theory, many have voiced concerns that subsidies do not encourage hiring or retention of workers in practice. Instead, employers may only apply for wage subsidies for workers whom they would have hired anyway, and so the subsidies have minimal effects on the employment of targeted populations (Bartik, 2001; Hamersma, 2005). Others have expressed concern that hiring subsidies may lead to "churning" of workers, in which subsidized employees are fired after the subsidy is exhausted (Government Accountability Office, 2001), or that targeted wage subsidies may stigmatize eligible workers and make them less likely to be hired (Burtless, 1985).

To estimate the effect of WOTC on labor market outcomes, we use administrative data from over 13 million individuals in the state of Wisconsin over the past two decades. Our data links quarterly employment and earnings records to a host of social program data, including SNAP, Temporary Aid to Needy Families (TANF), unemployment insurance, and criminal justice records. We additionally link this to individual-level records of over 800,000 application and certifications for WOTC subsidies, which allows us to characterize which types of firms and workers receive WOTC benefits.

Our empirical focus is SNAP recipients, which comprise approximately two-thirds of all workers whose wages are subsidized by WOTC (Department of Labor, 2024).[1] We use four complementary empirical approaches to establish the effects of WOTC on labor market outcomes. Each approach provides insights into different elements of the program, but the results are remarkably similar, reinforcing the robustness of the findings. The first strategy is based on a major WOTC expansion among SNAP recipients in 2007. Prior to 2007, SNAP recipients aged 18-24 had been eligible for wage subsidies, while in 2007, eligibility was expanded to ages 18-39. We use a difference-in-differences approach, measuring how labor market outcomes change for SNAP recipients just below the new age threshold of 40 (newly eligible) as compared to those just above (ineligible). We also examine how outcomes change for those just above age 25 (newly eligible) as compared to those just below age 25 (always eligible).

Second, we exploit these sharp age-based eligibility criteria for WOTC among SNAP recipients at age 25 prior to 2007 and age 40 beginning in 2007. Using a regression discon-

[1]Other categories include TANF recipients, Supplemental Security Income (SSI) recipients, individuals with a felony conviction, and the long-term unemployed. SNAP is by far the largest target group: the next largest group (long-term unemployed) comprised only 6.8% of WOTC certified hires in 2023 (Department of Labor, 2024).

tinuity design around those age thresholds, we test how labor market outcomes differ for those who are just eligible for the subsidies as compared to those who are just ineligible. This approach provides a larger sample for exploring potentially heterogeneous effects of eligibility and allows us to estimate the consequences of WOTC eligibility throughout the sample period rather than just around 2007.

Third, Wisconsin instituted electronic filing of WOTC applications in June 2017. This decreased the (implicit) cost to firms of applying for WOTC subsidies and led to a large increase in WOTC utilization. We examine the effect of e-filing on labor market outcomes by using the same WOTC-eligible and ineligible age groups among SNAP recipients as in the first empirical approach (around age 40), but with a difference-in-differences approach around June 2017. This approach provides a conceptually different treatment effect estimate, as it evaluates the effect of a policy change that lowers the cost of firm access to wage subsidies rather than individual-level eligibility.

Across all three approaches, we find no evidence that WOTC increases employment of eligible workers. These estimates are extremely precise: for example, using a precision weighted average of three approaches, the upper bound of the 95% confidence interval around the (null) estimate on hiring is 0.2 percentage points. These results hold even when we focus on employment at firms that are the largest users of WOTC. This implies that WOTC eligibility is not influencing firm hiring decisions, and so firms are receiving subsidies for workers whom they would have hired in the absence of WOTC, a phenomenon known as "windfall wastage" (Bartik, 2001). To quantify the extent of windfall wastage, we revise each specification so that the outcome is whether the worker is hired into a WOTC-subsidized job; this measures the effective difference between the eligible and ineligible group in program uptake. We divide the bounds on our hiring estimates by these uptake estimates to establish bounds on the fraction of WOTC-subsidized hires whose hiring could be attributable to WOTC. Our estimate of this ratio implies that 96% of WOTC-subsidized hires are windfall wastage, and we cannot reject the null that all WOTC hires are windfall wastage.

Even if hiring decisions are unchanged by wage subsidies, they could still benefit targeted workers if passed through in the form of higher wages or if employers retain workers for longer to receive subsidies. Either of these channels would imply that eligibility conveys higher earnings. Using the same empirical approaches, but instead using earnings as the outcome variable, we find no evidence of this. These null estimates are again quite precise, where the upper bound of the estimated confidence interval implies that we can rule out positive effects on quarterly earnings as little as $20 per quarter.

Proponents of wage subsidies suggest that wage subsidies could reduce usage of public benefit programs and discourage engagement in criminal activity. Indeed, WOTC lobbyists claim that these effects are meaningful and results in WOTC creating a large net savings for the government (Cappelli, 2012). Using administrative data on these outcomes and the same empirical approaches, we find no evidence of WOTC eligibility affecting the dollar value of public benefits received or any of a broad set of outcomes related to criminal activity.

Our fourth empirical approach analyzes the data at the firm-level rather than worker-level. We exploit the staggered adoption of WOTC by firms over time in a triple-differences design, examining how firm hiring practices changes around the time of WOTC adoption relative to firms that have not yet adopted WOTC. In particular, we focus on changes in the difference between firm hiring of SNAP recipients who are eligible for WOTC (i.e., just below age 40) relative to SNAP recipients who are just ineligible (i.e., just above age 40). After adoption, there is an immediate and persistent spike in firm receipt of WOTC subsidies for eligible hires relative to ineligible ones, meaning that the firm faces a lower wage bill when hiring a WOTC-eligible worker. However, as in the previous three analyses, we find that they do not respond to this incentive: they are no more likely to hire eligible workers relative to ineligible workers.[2]

While we find no effect on hiring, an oft-cited concern with capped wage subsidy programs is that firms may fire subsidized employees after the subsidy is exhausted to hire a new batch of subsidized workers, a phenomenon labeled as "churning"(Government Accountability Office, 2001). If this were the case, hiring subsidies may reduce the availability of stable jobs in low-wage labor markets. We investigate this possibility using discontinuities in the WOTC subsidy schedule, checking for bunching or kinks in termination after subsidy exhaustion or other discontinuous changes in subsidy amount. We find no evidence of churning, including

[2]We also find no evidence of aggregate effects at the firm level. Using a strategy that compares firms that are more or less "exposed" to the 2007 policy change, more exposed firms are no more likely to hire more workers on aggregate or increase wages than less exposed firms.

when we focus on particular categories of firms that are more attentive to WOTC or within particular time periods.

Put together, our results imply that hiring subsidies through WOTC are a pure transfer to firms. Moreover, we find that these transfers are heavily concentrated, with half of WOTC subsidies in Wisconsin going to just 48 firms, far in excess of these firms' share of employment (9% of all hires). Thus, in sharp contrast to the stated objectives of aiding disadvantaged workers in finding jobs, the wage subsidies appear to instead accrue to a small set of firms.

The final section of the paper examines the puzzle of why firm hiring practices do not respond to this relatively sizable subsidy. We first test and rule out conventional explanations related to standard models of labor supply and demand or the negative signal value of WOTC eligibility. We then collect data on firm hiring processes for a representative sample of WOTC-using firms. These data provide a clear and stark explanation: hiring decisionmakers almost never know who is WOTC eligible. Even among these firms that receive WOTC subsidies, less than 15% of initial job applications collect information that would allow firms to determine an applicant's WOTC eligibility, and even among those that do, the information is often explicitly siloed from individuals making hiring decisions. To understand why firms do not integrate WOTC status into hiring, we conduct a separate survey of individuals with hiring experience at firms that receive WOTC subsidies. We find that firms are concerned that preferring WOTC-eligible applicants could risk lawsuits over hiring discrimination, and that explicit screening may discourage applications due to imperfect information on the applicant side. These channels are typically absent from models of hiring and job search, but appear to generate important deviations from predictions of standard theory.

This paper contributes to three main literatures. First, our results bring new estimates to the literature on hiring and wage subsidy policies in the United States. An older literature studied early US federal wage subsidy programs such as the New Jobs Tax Credit (e.g., Bishop, 1981; Perloff and Wachter, 1979) and the Targeted Job Tax Credit (e.g., Bishop and Montgomery, 1993; Burtless, 1985; Hollenbeck and Wilke, 1991; Katz, 1996; Lorenz, 1995), but Katz (1996) suggests concerns with the empirical designs of many of these studies, concluding that there "exists much uncertainty [about the efficacy of wage subsidies]." More recent papers have studied smaller state tax credits (e.g., Bartik, 2001; Bartik and Erickcek, 2014; Neumark and Grijalva, 2017), test how employers respond to wage subsidies in a survey or experimental context (Bushway and Pickett, 2024; Cullen et al., 2023; Hunt et al., 2018), or examine subsidies packaged with other interventions like training (e.g., Bell and Orr (1994); Gelber et al. (2016)).[3] We provide estimates for a much larger federal program, and for wage subsidies that are targeted at individuals with disadvantages in the labor market.

Second, our paper contributes to the literature in labor economics highlighting the consequences of frictions typically absent in classic models of labor markets. A recent literature has shown, for example, that equity pay constraints can explain the lack of within-firm pass-through of payroll taxes (Saez et al., 2019) and worker-side information frictions and misperceptions in job negotiations can impact wage gaps (Cullen and Pakzad-Hurson, 2023; Jäger et al., 2024; Roussille, 2024). Our results show that firm-side information frictions are also important in determining labor market outcomes, and even that these frictions can be self-imposed in response to legal institutions.

Third, we contribute to a smaller literature that has studied WOTC specifically. These papers either focus on an earlier time period at which point the program was relatively small and focused on TANF recipients (Hamersma, 2008) or are descriptive in nature (Gunderson and Hotchkiss, 2007; Hamersma and Heinrich, 2008; Hunt et al., 2018). The lack of academic consensus on WOTC is well-illustrated by a February 2023 release by the Department of Labor containing a list of open questions about the program such as "does the tax credit influence employer hiring decisions?" (Department of Labor, 2023a). Our results expand and often differ from the findings in existing work, providing an overall more negative assessment of the program.[4] While our empirical results focus on SNAP recipients (the largest target

[3]Outside of the US, Cahuc et al. (2019) finds that hiring credits for French firms during the Great Recession increased employment without affecting wages, Boockmann et al. (2012) finds that hiring subsidies for older workers in Germany only increase the hiring of some types of workers, and studies in other European countries have found modest or no effects (e.g., Huttunen et al., 2013; Schünemann et al., 2015). Fenizia et al. (2025) finds a similarly precise null effect of wage subsidies for hiring apprentices in Italy, but shows that this results from a different underlying mechanism than this paper (inelastic demand for apprentices among firms). There have also been a number of recent evaluations of wage subsidy programs in developing countries using randomized controlled trials, which typically find short-term employment gains without longer-run effects (e.g., Abel et al., 2022; Galasso et al., 2004; Groh et al., 2016).

[4]The only other paper that attempts to causally identify the labor market consequences of WOTC, Hamersma (2008), estimates that WOTC generates a meaningful (5.9 percentage point) increase in short-run employment, but not in long-run employment. Our findings may differ due to sample, timing, or empirical strategy: Hamersma (2008) uses data from approximately 3,500 TANF recipients between 1997-2002 and compares labor market outcomes of individuals who stopped receiving TANF benefits after 6-8 months (ineligible for WOTC) to individuals who received TANF for 9-18 months (eligible).

group) and thus we cannot definitively speak to the effects of WOTC on other target groups, the informational constraints that we identify also apply to most of those groups. The Work Opportunity Tax Credit will be up for renewal by Congress in December 2025, including pending legislation that proposes to expand the scope of the program (Congress, 2023). Overall, our results imply that either reforms are needed to make WOTC actually benefit targeted groups or that the billions spent annually on WOTC could be more impactful if allocated to other programs.

2.3 Institutional Background

Active labor market policies, such as job search assistance, job training, and job subsidies, are a large part of government spending in most developed countries: OECD countries spend an average of 0.58% of GDP on these policies annually (Crépon and Van Den Berg, 2016). While a large literature on job search assistance and job training programs suggests that the former can sometimes be helpful and the latter is often not, there is much less evidence on employment subsidies, particularly in the US context.[5]

The US has a long history of experimenting with wage subsidy programs, including Job Opportunities in the Business Sector in the late 1960s, the WIN tax credit in the 1970s for AFDC recipients, and the New Jobs Tax Credit from 1977-1978 (see, e.g., Bishop, 1981; Perloff and Wachter, 1979). The most extensive of these programs was the Targeted Job Tax Credit, which was in effect between 1979 to 1994 and primarily subsidized the employment of beneficiaries of public assistance. This program was eventually phased out after criticism over implementation issues, namely "employer windfalls for hiring employees that they would have hired anyway and too many credit-eligible employees leaving their jobs before receiving much work experience" (Government Accountability Office, 2002). This paved the way for reforms to the design of federal targeted wage subsidies that resulted in the Work Opportunity Tax Credit.

[5]See Card et al. (2010), Crépon and Van Den Berg (2016) and Card et al. (2018) for meta-analyses and summaries on job search assistance and job training programs.

2.3.1 The Work Opportunity Tax Credit

The Work Opportunity Tax Credit was established by Congress on August 20, 1996 by P.L 104-188. The current WOTC differs from past wage subsidy programs in both design and target populations. As discussed in a 2002 GAO report on WOTC, the most important change was introducing a requirement that employers confirm applicants' WOTC eligibility prior to making hiring decisions to avoid windfall wastage (Government Accountability Office, 2002): on IRS form 8850, which firms submit when filing for WOTC, both the hire and firm must sign to confirm that "under penalties of perjury, I declare that the applicant provided the information on this form [indicating eligibility for WOTC] on or before the day a job was offered to the applicant." This contrasts with the Targeted Job Tax Credit, under which employers "may claim the tax credit even though they have made no specific effort to recruit, hire, or retain workers targeted by the program" (Government Accountability Office, 1991). There were also changes to the minimum employment length requirements to encourage longer stints of employment, as well as reforms to eligibility target groups to focus on the neediest job seekers (Government Accountability Office, 2002). Finally, as discussed more below, the number of subsidized workers has expanded substantially: the Targeted Job Tax Credit cost around $300 million in 1994 (U.S. Department of Labor Office of Inspector General, 1994), while we estimate that WOTC currently costs over $2 billion annually. These changes have the potential to make WOTC significantly more effective than its predecessors, but it is an empirical question whether this potential manifests in actual gains.

The main groups of individuals eligible for WOTC are those: (1) aged 18 to 39 and whose family has received SNAP benefits for at least 3 of the 5 months prior to hiring; (2) whose family has received assistance under TANF for at least 9 of the 18 prior months; (3) convicted of a felony and hired within one year of their conviction or release from incarceration; (4) who are long-term TANF recipients, including those who exhausted their TANF benefits in the past two years; (5) who have received SSI over the past three months;[6] or (6) unemployed for at least 27 consecutive weeks and who received unemployment compensation during at

[6]See appendix C.1.3 of Aizawa et al. (2025) for an analysis using the Health and Retirement Study on how WOTC affects employment among individuals with disabilities. While they find no effect of WOTC on employment outcomes, they note that their estimates are noisy.

least part of the unemployment period.[7] SNAP recipients are by far the largest category, making up around two-thirds of all WOTC certifications in recent years (Figure 2.1). These populations typically have low rates of labor force participation: for example, only 52% of non-disabled, working-age SNAP participants are employed in a typical month (Keith-Jennings and Chaudhry, 2018).

All firms are eligible to apply for WOTC and there is no limit on the number of hires for whom they can receive subsidies, but they can only receive the subsidy for a particular worker once.[8] However, even if another firm has received WOTC for a particular worker in the past, the worker's current employer can still claim the full subsidy amount as long as the worker satisfies the eligibility criteria at the time of hiring. Employers need only submit two one-page forms for each worker within 28 days of hiring to the relevant state agency (plus a one-page form to the Internal Revenue Service at the end of the year aggregating all individual claims) to claim the tax credit.[9] This agency verifies that the worker is eligible and certifies or denies the application.[10] Using data from three states for which we observe data on the exact date of filing an application and certification for WOTC applications, we find that the median time between application and certification is 119 days, with a standard deviation of 115 days.[11]

When applications are approved, the firm receives tax credits for wages over the worker's first 12 months of employment in a job. The firm can claim 25% of wages for workers who worked between 120-400 hours and 40% for employees who worked at least 400 hours, up to a limit of $6,000 of earnings (or $2,400 in tax credits) per worker for most categories of

[7] A few other categories are also eligible, such as qualified veterans, residents of "designated communities", and employees of summer youth programs, but make up a small fraction of all WOTC subsidized hires

[8] Taxable employers can claim WOTC against income taxes, while eligible tax-exempt employers can only claim WOTC against payroll taxes on wages paid to qualified veterans rather than the full set of eligible individuals listed above (Internal Revenue Service, 2024).

[9] The form also requires applicants to supply some evidence of eligibility; for example, for those eligible based on TANF or SNAP receipt, applicants need to provide their case number identifier and/or a signed statement from an Authorized Individual with a description of months of benefits that were received.

[10] While there is an option for workers to get "pre-certified", so that they can present their eligibility to potential employers and avoid this application process, this accounts for a minuscule fraction of certifications in our data (0.65%), almost all of which are eligible due to a felony conviction.

[11] These data are for Arizona, Tennessee, and Rhode Island, and come from FOIA requests (Corwin, 2022). Using coarser (quarterly) data from the Department of Labor that records the number of WOTC applications, certifications, and rejections in a quarter, we estimate the number of quarters in which applications are processed in each state. The processing speed of these three states is close to the national average, so the estimated processing time for those three states is likely to generalize.

workers (Appendix Figure B1).[12] These tax credits are non-refundable and so require tax liability to be claimed, but they can be carried backward up to one tax year and forward up to 20 tax years. The firm also continues to receive the subsidy even if the worker does not remain in the eligible category after hiring.

WOTC primarily subsidizes low-wage employment. Appendix Figure B2 plots a histogram of the starting hourly wage for WOTC workers, using administrative micro-data on all WOTC applications in Wisconsin between 2005 and 2020. Most jobs are at or just above the minimum wage, which was equal to $7.25 in Wisconsin over this period.[13] Combining this data with administrative quarterly earnings data, we find that for each certification, firms receive an average of $1098 in subsidies. 34% of certified workers work less than the minimum number of hours for eligibility for WOTC (120 hours), 24% work at a 25% wage subsidy rate (120-400 hours), and 42% work enough hours to qualify for the 40% subsidy rate, with 31% of all certifications yielding the maximum subsidy amount of $2400.[14] As these numbers indicate, these jobs are typically short-term, with a median estimated hours worked of 265 hours and only 23% of jobs lasting longer than three quarters.

2.3.2 History of the Work Opportunity Tax Credit

In its early years, WOTC subsidized the wages of approximately 400,000 individuals annually, where the majority were recipients of TANF. Over time, eligibility criteria have expanded to include additional groups, such as SNAP recipients between the ages of 25 to 39 (in 2007) and the long-term unemployed (in 2015). As a result of these extensions and greater awareness of the program, the number of workers whose wages are subsidized by WOTC has grown sixfold, reaching 2.5 million in 2022 (Department of Labor, 2023b). Panel (a) of Figure 2.1 shows that most of this growth came between 2008 and 2015, and is due to certifications

[12] The limit is $9,600 in tax credits for disabled veterans who meet other criteria, but this is a very small fraction of all WOTC hires.

[13] Data from administrative records in seven other states between 2018 and 2020 (Corwin, 2022) suggest that this is also true outside of Wisconsin.

[14] To calculate total hours worked for the employer, we take the starting hourly wage reported in the WOTC application and divide total earnings for the worker from that employer in this work stint. This could be incorrect if workers receive a raise in their hourly wage, but that is quite unlikely since 400 hours is still only 10 weeks of full-time work and so workers are unlikely to receive a raise so quickly.

for WOTC through SNAP.[15] Consistent with this, panel (c) of Figure B3 uses data from the Current Population Survey to estimate the fraction of individuals below the poverty line who would be eligible for WOTC if hired into a new job. This number expands sharply after an expansion in WOTC, and after 2010, ranges from 36 to 47% of all adults between 18 to 40 who are below the poverty line. Patterns in Wisconsin are quite similar to those nationally, suggesting that even though our data is from one state, it will well approximate national patterns.

To further decompose the reasons for this massive growth, we leverage individual-level data on all WOTC applications in the state of Wisconsin between 2005 and 2019, a dataset described in greater detail in Section 2.4. We match this to employer-employee earnings records, as well as TANF, SNAP, and criminal justice records to measure which new hires are potentially eligible for WOTC subsidies. The expansion could be the result of: (i) an increased number of job applicants who are eligible for WOTC; (ii) firms becoming more likely to apply for WOTC subsidies for a given hire, such as if screening technology improves; or (iii) greater success in the application process, such as a higher fraction of applications being accepted.

Panel (a) of Appendix Figure B3 plots the total number of new hires who are potentially eligible for WOTC over time as a result of the three eligibility criteria that are our observable in our data – SNAP receipt, TANF receipt, and having a recent felony conviction – which account for over three-quarters of WOTC subsidized hires. We then match these employment stints to WOTC micro-data to check what fraction of eligible hires have an associated WOTC application, which we plot in panel (b). Panel (b) also plots the fraction of all applications that result in a certification. Over the main period of expansion from 2008 to 2015, there is a tripling in the number of hires who qualify for WOTC. The fraction of eligible hires for whom firms apply for WOTC also increased meaningfully from 11.2% in 2007 to 17.1% in 2015, though this is less important than the increase in eligibility.[16] Finally, application

[15]Federal law places limits on SNAP eligibility for able-bodied adults without dependents (3 months in a 36 month period). Wisconsin had these requirements waived between 2002 and 2015 due to high unemployment, and we do not see evidence of differences in the effects of WOTC during the waiver/non-waiver periods.

[16]Note that this is the fraction of *all* hires who would be WOTC eligible, where many of these hires are at firms that do not engage with the WOTC program. Among the set of firms that engage with WOTC, applications are filed for around 40% of WOTC-eligible hires, without much change in this rate over the study period.

success rates are fairly steady, and so do not explain the change. We thus conclude that the growth of WOTC is mostly due to an increase in the number of hires who are eligible, largely due to the expansion of the SNAP program.

2.4 Data

The major difficulty with studying WOTC is that the necessary administrative data on eligibility and outcomes are typically siloed within many different state government agencies. We overcome this challenge by using linked administrative data from the state of Wisconsin through the Institute for Research on Poverty (IRP) at the University of Wisconsin-Madison, which contains information on over 13 million individuals. We use linked data from quarterly earnings records, SNAP benefits, TANF benefits, criminal justice records, unemployment insurance benefits, WOTC application and certification data, and a data core that compiles individual demographic information from various state databases. Wisconsin is a good state in which to study WOTC because it is fairly typical in its implementation of the program: for example, Appendix Figure B4 shows that per capita utilization is almost exactly the national average.

Quarterly earnings records: We use quarterly earnings records between 1998 and 2020 that are collected for Wisconsin's Unemployment Insurance (UI) program from the Department of Workforce Development.[17] These records include quarterly earnings for each worker from each of their employers, along with worker and employer identifiers.

SNAP records: We use monthly FoodShare (the name of Wisconsin's SNAP program) benefit records between 1998 and 2020 from the Department of Health Services. These records include benefit amounts for the household and individual identifiers of each eligible member of the household.

[17] While this data excludes some workers, including those who work in a neighboring state, those who have moved out of state, Federal employees, and self-employed workers and independent contractors, these are a small fraction of total employment.

TANF records: We use monthly TANF benefit records between 2003 and 2019 from the Department of Children and Families. These records include benefit amounts and individual-level identifiers.

Criminal justice records: These data, which cover the years 2004 to 2020, consist of court records, state Department of Corrections data on entry/exit from prisons, and data on entry and exit from Milwaukee jails. For the purposes of this study, we combine these to observe whether an individual has been convicted of a misdemeanor or felony in a given quarter.

Work Opportunity Tax Credit application and certification data: We use WOTC certification data between 2005 and 2020 from the Department of Workforce Development. These records include the date of the WOTC application, whether the application was certified or denied, and starting in 2009, the target group under which the certification qualified. It also includes individual and firm identifiers which allows us to link the WOTC certification data to the other data sources.

Demographic data: For individuals who receive a social service such as SNAP, we also observe demographic information such as age, gender, and race.

Table 2.1 shows summary statistics for (1) individuals for whom a firm files a WOTC application; (2) individuals for whom the firm successfully receives a WOTC certification; and (3) the sample of individuals who are between the ages of 18-39 and receive SNAP benefits. Among the WOTC applicants and certified individuals, our data consists of 797,411 applications and 426,498 certifications across over 400,000 individuals. The average age for both applicants and certifications is around 29 years, less than half the applicants are male, and over a third are Black. Over three-quarters of successful WOTC certifications are for individuals who are SNAP beneficiaries, while only 10% of them receive TANF benefits. Median quarterly earnings of WOTC-certified individuals is approximately $1,800 per quarter. For the WOTC-certified job in particular, the job lasts on average around 3 quarters, with median total earnings of around $2,600 at an average starting wage of $9/hour.

Retail, admin/support, and food/accommodation are the largest industries, and the average WOTC subsidy is a little over $1,000.

The SNAP sample of individuals, which consists of almost 900,000 individuals, is comparable demographically to individuals who receive WOTC certifications. Job stints are also comparable in terms of median total earnings and tenure with a firm, though the jobs are much less likely to be in retail.

2.5 Effects of WOTC eligibility on individual outcomes

In this section, we use the linked administrative data to estimate the effects of WOTC eligibility on labor market outcomes – employment, being hired into a new jbo, and earnings – as well as spillovers to usage of social assistance and criminal activity. We use three separate identification strategies based on individual-level data: first, we exploit a 2007 expansion in eligibility among the SNAP recipient target group to older ages in a difference-in-differences design; second, we exploit sharp age cut-offs in WOTC eligibility in a regression discontinuity design; and third, we leverage a 2017 change in the application process that implicitly decreased the cost of applying for WOTC.

2.5.1 SNAP expansion

We first examine a large expansion of eligibility among SNAP recipients. SNAP recipients are the largest target group for WOTC, comprising around two-thirds of WOTC certifications in recent years. Prior to 2007, individuals who were between the ages of 18-24 and had received SNAP benefits in three of the previous five months were eligible for WOTC. In 2007, this target group was expanded to include SNAP recipients between the ages of 25 to 39. Figure 2.2, which graphs total WOTC certifications by application date in our data, shows that total certifications increased from an average of 3,000 per quarter for hires in 2005-2006 to an average of over 5,000 per quarter for hires in 2007-2008 after the SNAP expansion.[18]

To study the impact of this expansion, we construct a sample of individuals who have

[18]Our data on WOTC certifications only includes the reason for eligibility after 2009. Thus we focus on total certifications rather than only certifications due to SNAP eligibility here. A similar uptick is observed for SNAP-specific certifications using aggregate national data in Figure 2.1.

received SNAP benefits in three of the last five months. We then conduct two analyses (and also stack the analyses): in the first, we compare individuals in the sample who are slightly older than age 25 (eligible for WOTC after 2007) before and after the policy change to individuals slightly below age 25 (always eligible for WOTC). Analogously, we compare individuals who are slightly younger than 40 (eligible for WOTC after 2007) before and after the policy change to individuals above age 40 (ineligible for WOTC). This design works well because to the best of our knowledge, these age cutoffs were not used as a cutoff for other programs experiencing reforms around this time. We focus on the sample period between July 2005 through December 2008 and run event studies and difference-in-differences regressions.

The event study specification for the age 40 cutoff is:

$$Y_{i,t} = \sum_{s=2005q3, s \neq 2006q4}^{2008q4} \beta_s 1\left\{age_{i,t} < 40\right\} + \eta_t + \gamma_a + \varepsilon_{i,t} \tag{2.1}$$

where $Y_{i,t}$ is an outcome for individual i measured in quarter t, η_t are year-quarter fixed effects, and γ_a are age (in months) fixed effects. $age_{i,t}$ measures individual i's age at the start of quarter t, and the indicator is equal to one for individuals whose age is less than 40 (the treatment group) and zero for individuals whose age is 40 or over (the control group). The treatment group consists of individuals with ages between 37.75 years and 39.75 years at the start of a quarter: we select this cutoff as those individuals will be eligible for WOTC for the entirety of quarter t after 2007. Our control group consists of individuals aged between 42 and 44. We exclude individuals aged 40-41 so that the control sample consists fully of those who were not previously eligible for WOTC over the eight quarters after WOTC eligibility expansion in the first quarter of 2007. For example, an individual who turns 40 in the second quarter of 2007 would have been 39 (and thus treated) in the first quarter of 2007.

We also estimate the corresponding difference-in-differences specification that collapses the β_s coefficients into one coefficient post-policy change relative to all time periods prior to the policy change:

$$Y_{i,t} = \beta^{DD} \times 1\{age_{i,t} < 40\} \times 1\{t \geq 2007q1\} + \eta_t^{DD} + \gamma_a^{DD} + \epsilon_{i,t} \tag{2.2}$$

For both the event study and difference-in-differences specifications, we cluster standard errors at the individual level.

We also run analogous specifications for the age-25 cutoff, where the treatment group consists of individuals aged between 25 and 27 years at the start of a quarter and the control group consists of individuals aged between 22.75 and 24.75; for similar reasons as the age-40 cutoff, we select the 24.75 cutoff to ensure that all individuals in the control group remain eligible for WOTC for the entirety of quarter t prior to 2007.

We begin by examining three main labor market outcomes. Our first outcome is new hires, defined as an indicator for whether the individual had positive earnings in quarter t in a firm at which they did not have positive earnings in the prior quarter. Our second outcome is employment, defined as an indicator for whether individual i had positive earnings in quarter t. An advantage of using the new hires outcome over employment is that it measures flows into employment, the target of the policy, and excludes workers who are persistently employed. Our third outcome is earnings in quarter t. This measure captures tenure in the job and wage rate, both of which might be affected by WOTC. For each outcome, we examine effects for the full sample as well as the sample of individuals who were unemployed at the start of period t, and so for whom WOTC may be particularly relevant.

Columns (1)-(3) of Table 2.2 run the difference-in-differences regressions, with Panel (a) showing the stacked effects and Panels (b) and (c) showing the effects separately for the age-25 and age-40 cutoffs. Panels (a)-(c) of Figure 2.3 report event study coefficients from Equation (2.1) for these three labor market outcomes, stacking the two age cutoffs (see Appendix Figures B5 and B6 for the age-25 and age-40 cutoffs, respectively). For all three outcomes, there is no evidence of a pre-trend in outcomes prior to the policy change, and also no detectable effect of the expansion on the outcomes in the eight quarters after the expansion. The confidence intervals on β_{DD} are tight: for example, the upper bound of the 95% confidence interval on being hired would be a 0.25 percentage point increase (or 1.6%

of the base new hiring rate of 15.5% for this population) and a $66 increase in quarterly earnings (or 3.4% of the $1,934 in average quarterly earnings for this population).[19]

These results suggest that WOTC eligibility does not meaningfully improve labor market outcomes among the SNAP target group. This is not driven, however, by a lack of use of WOTC. Panel (d) of Figure 2.3 and Column (4) of Table 2.2 show event study coefficients and difference-in-differences estimates running the same specification with the outcome as being newly hired into a WOTC-certified job in the quarter. There is an immediate increase in the likelihood of being hired into a WOTC-certified jobs for the newly eligible group relative to the never eligible (control) group of 1.1 percentage points.

Finally, lobbyists for WOTC frequently claim that the program results in net savings for the government by allowing people to obtain jobs that pull them off of public assistance and out of criminal activity (Cappelli, 2012). Our precise null estimates on labor market outcomes suggest that this is unlikely to be true, but given the policy importance of this question, we also examine those outcomes. We construct a measure of social assistance as the dollar amount of benefits the individual received from SNAP, TANF, or unemployment insurance in time t. For criminal activity, we construct an indicator equal to one if the individual has a criminal conviction at time t. Panels (e) and (f) of Figure 2.3 and columns (5) and (6) of Table 2.2 show similar null effects for these two outcomes. It is thus unlikely to be true that these wage subsidies generate any savings for the government.

These results are robust to sub-samples and alternative specifications. Appendix Figures B2 and B3 shows null employment and hiring results by race, gender, past conviction status of individuals as well as outcomes by firm type (top 50 WOTC-hiring firm, firms with high success rates on WOTC applications, staffing firms, and firms in highly competitive industries). Appendix Figure B7 and Appendix Table B1 show that our results are similar when we narrow our sample to individuals who enter quarter t unemployed. This sample restriction provides additional precision by focusing on a population that is more likely to be seeking a job. Another concern is that WOTC could have spillovers across workers, whereby

[19]In a highly related paper, Cullen et al. (2023) measure how wage subsidies affect employer willingness to hire a worker with a criminal record, experimentally varying the subsidy amount from 0% to 100%. Although they study a different and more heavily stigmatized population (workers with a criminal record), our results turn out to be fairly consistent with theirs: they find no response to wage subsidies of up to 25%, which is close to the effective subsidy rate for WOTC calculated in Section 2.3.

subsidized employment for workers in the treatment group may displace employment in the control group. However, such spillovers would *increase* differences in employment outcomes between treatment and control rather than bias estimates towards zero. Thus the fact that we find null effects would tend to dispel this concern.[20] The following sections provide further evidence of robustness with alternative empirical approaches.

2.5.2 Age-based discontinuities for SNAP

As a second method of evaluating the effect of WOTC, we exploit the sharp age-based discontinuities in eligibility for WOTC among SNAP recipients. Prior to 2007, SNAP recipients just below age 25 were eligible for WOTC, whereas those just above were not; after 2007, there is a similar discontinuity around age 40. An advantage of this approach is that we will be able to estimate the consequences of WOTC eligibility throughout the sample period rather than at just two points in time (SNAP expansion in 2007; introduction of electronic filing in 2017). It also generates a very large sample, providing excellent statistical power for our tests; we use this to explore potential heterogeneity in treatment effects.

We combine two sharp regression discontinuity designs: around age 25 for the period 1998 to 2006 and around age 40 for the period 2007 to 2019. Our estimates are based on the following specification:

$$Y_{i,t} = \gamma_0 + \gamma_1 1\left\{age_{i,t} \geq A_t\right\} + f_1\left(age_{i,t}\right) + f_2\left(age_{i,t}\right) 1\left\{age_{i,t} \geq A_t\right\} + \epsilon_{i,t} \tag{2.3}$$

where $Y_{i,t}$ is the labor market outcome for individual i in quarter t, $age_{i,t}$ measures their age at the start of quarter t, and A_t is equal to 25 for years $<$ 2007 and 40 for years $\geq$ 2007. The coefficient γ_1 captures the (negative) treatment effect, while the functions f_1 and f_2 reflect a continuous but potentially non-parametric relationship between the running variable and the outcome.[21] Using the data-driven approach of Calonico et al. (2014) and *rdrobust* command of Calonico et al. (2017), we apply a triangular weighting kernel in distance from

[20] Furthermore, substitution is unlikely to come purely from the control group: for example, in the SNAP expansion research design, it could be that employment of the treatment group (SNAP recipients aged 38-39) displaces someone not from the control group (SNAP recipients aged 42-43) but SNAP recipients aged above age 43 or non-SNAP recipients in any age range.

[21] Crossing the age threshold entails a change from treated to untreated, and so γ_1 is equal to the treatment effect multiplied by negative 1.

the threshold, calculate a MSE-optimal bandwidth with a linear polynomial estimated within the bandwidth on either side of the cutoff, and calculate heteroskedasticity-robust standard errors clustered at the individual level.[22]

As in the SNAP expansion analysis, we restrict the sample to individuals who have received SNAP for three of the five months preceding the start of a given quarter, i.e., the condition for WOTC eligibility through SNAP. We also drop the quarters in which $age_{i,t}$ crosses the relevant threshold (i.e., individuals whose age at the beginning of a quarter is (24.75,25) or (39.75,40)), as those are difficult to interpret since individuals are partially treated.

This approach has a number of key identifying assumptions. The first is that the running variable, age, is not manipulable around the threshold, such as if individuals lie about their age when applying for SNAP benefits. We test for continuity of the density of the age distribution within our sample around the cut-offs (Cattaneo et al., 2018), as seen visually in Appendix Figure B8: the p-value is 0.27, indicating a lack of manipulation. This also rules out the concern that individuals may be more likely to apply for SNAP benefits when below the relevant age threshold in order to make them eligible for WOTC (and possibly increase their desirability to employers), where this would predict greater density to the left of the threshold.[23]

The second concern is continuity across the age threshold of covariates that may be related to the outcomes, i.e., the composition of sample individuals changes discontinuously around this age threshold. This might be true if WOTC increases employment levels, reducing SNAP utilization and thus changing the composition of SNAP recipients above and below the threshold.[24] We will directly test how employment varies across the threshold below. Here, we examine whether any of the covariates we can observe change discontinuously around the threshold, which could have other root causes, such as other policies using this age cut-

[22]The MSE-optimal bandwidth and polynomial will vary for each outcome Y, following Calonico et al. (2018).

[23]We also use the full WADC data on Wisconsin residents to test this directly, where appendix table B5 finds that likelihood of receiving SNAP benefits does not vary around the age threshold.

[24]It is not clear whether changes in employment would change SNAP utilization since WOTC jobs are typically quite low wage. The SNAP income limit for a family of 3 in 2010 was \$36,620 annually, which means that a household with two adults working full-time at the median WOTC wage in our sample (\$9 per hour) would still qualify for SNAP.

off. We focus on gender and race, as well as employment and earnings in previous periods; the labor market variables are particularly good tests, as those are strongly correlated with labor market outcomes in time t, and so pre-existing discontinuities in those variables are most likely to generate bias. Pooling the pre- and post-2007 data, Appendix Table B4 and Appendix Figure B9 find no evidence of any of these variables varying discontinuously across the age threshold.

We first examine how WOTC eligibility affects labor market outcomes in Table 2.3 and Figure 2.4. This analysis combines both the age 25 (pre-2007) and age 40 (post-2007) discontinuities by defining the running variable as distance from the relevant age cut-off; Appendix Tables B6 and B7 report the results separately for the pre- and post-2007 periods. Column (1) of Panel (a) of Table 2.3 shows that hiring does not respond to WOTC eligibility, while column (3) finds no effect on earnings. For the RD analysis, we prefer looking at new hires rather than employment because it isolates flows into jobs, removing relatively persistent employment that we would not expect to respond to WOTC as well as holdover of jobs across the age threshold. Nonetheless, we also examine employment (column 2) and find no effect there. These null estimates are even more precise than the estimates based on SNAP expansion – with standard errors that are two to three times smaller – reflecting the greater sample size in the RD analysis.

These results broadly support the findings of the previous section that WOTC eligibility does not meaningfully improve labor market outcomes among targeted workers. To put these estimates in context, we estimate the extent to which individuals above and below the threshold differ in their likelihood of being in a WOTC subsidized job. Column (4) of Table 2.3 and panel (d) of Figure 2.4 show that individuals below the discontinuity are 0.97 percentage points more likely to be hired into a WOTC subsidized job.[25]

One concern with this analysis is that individuals may remain in a job for multiple quarters. An individual who receives a WOTC job at age 39 may still be employed in that job at age 40, which would tend to bias our estimates towards zero. In Panel (b) of Table 2.3 and Appendix Figure B10, we overcome this concern by restricting the sample to individuals

[25] Note that some individuals to the right of the threshold may still be hired into WOTC jobs by virtue of other eligibility criteria, such as a felony conviction or receipt of TANF.

who are unemployed at the start of the quarter t, and so would not experience job holdover. We find that all of the estimates are quite similar to those in Panel (a), so this does not appear to be a major concern. That is likely because most WOTC-subsidized jobs are very short-lived, as we will show in Section 2.6, reducing the extent of bias due to holdover.

Even though our estimates suggest that firms on average do not respond to WOTC eligibility in their hiring decisions, it could be that there is a subset of firms that do. We check for differential responses among four types of firms that may be more likely to respond to WOTC. First, we focus on employment at the 50 firms with the largest number of WOTC certifications over the sample period; these presumably may have the best WOTC screening procedures. Second, we refine the measure of attentiveness to WOTC by calculating the fraction of WOTC eligible hires for which each firm actually applies for WOTC. We then focus on employment in firms whose application rate is in the top 25% of WOTC-using firms, as these are the firms that appear to pay the most attention to WOTC. Third, employment services agencies, which provide temporary employees to businesses, are the biggest category of firms utilizing WOTC, and so may be most attentive to the policy. Fourth, firms in industries that are more competitive may be more responsive to policies like WOTC due to their lower margins. We focus on firms in the top 25% most competitive industries (four digit NAICS codes) based on data from the 2022 US Economic Census. In columns (1) to (4) of Appendix Table B9, we re-estimate the main RD specification, but redefine the outcome variable to be whether the individual is hired in each of these four categories of firms. Even among these firms where the treatment effect may plausibly be larger, we still find no response for either employment (panel a) or new hiring (panel b).

One concern with the approach in Section 2.5.1 is that it was based on an expansion in 2007, but the WOTC program has grown considerably since then. It could be that as part of this growth, HR departments have implemented screening for WOTC-eligible job applicants such that WOTC now has an effect on hiring. An advantage of our RD approach is that we can re-estimate the discontinuity at different points in time. The remaining columns of Appendix Table B9 subset the data to different periods of time: column (4) restricts to the pre-2007 period (with the age 25 cut-off), column (5) focuses on the period during which there was significant growth in WOTC (2007 to 2013), and column (6) examines a more recent period (2014-2019), in which number of WOTC certifications stabilized. However, we find no evidence of differential effects over time.

Another concern with the RD approach is that other policies may use the same age thresholds; for example, if there were labor market support policies targeted at individuals above age 40, this could bias our estimates downwards, and potentially explain the null effect. As a final robustness check, we conduct a placebo analysis examining how labor market outcomes vary across the age thresholds during years when the threshold were not relevant for WOTC: i.e., look at the age 40 discontinuity in the pre-2007 data and the age 25 discontinuity in the post-2007 data. If there were other policies targeting these age thresholds that also affected labor market outcomes, we should observe these effects in this placebo analysis. Appendix Table B10 finds no differences over the placebo age thresholds, consistent with the RD approach cleanly identifying the effect of WOTC.[26]

2.5.3 Effects of reducing WOTC application costs

The previous analyses examined the effect of expanding eligibility for wage subsidies under WOTC. This section analyzes a reduction in the cost to firms of filing for WOTC through the introduction of electronic filing (eWOTC), which occurred in June 2017. As shown in Panel (a) of Figure 2.5, this led to a substantial increase in the number of WOTC certifications: total WOTC certifications increased 50% from an average of 8,000 per quarter from 2015 to June 2017 to an average of 12,000 per quarter from July 2017 to 2019.[27] To check that this increase was not driven by aggregate changes to labor markets, Panel (b) uses Department of Labor data on quarterly state-level WOTC applications to compare WOTC uptake in Wisconsin to the neighboring state of Minnesota, which did not adopt e-filing at this time.[28] Although the two are similar prior to June 2017 in both levels and trends, certifications

[26]The coefficient on WOTC certification is just statistically significant at the 10% level for the post-2007 period ($p = 0.08$). This is entirely due to data from 2007 – i.e., the year that the new age threshold was put into place – where the coefficient is no longer statistically significant when the first two quarters of this year are removed. This is likely because it takes time to learn about the new age eligibility cutoff, where some firms initially continued to follow an age 25 cut-off for submitting WOTC applications before learning of the new threshold.

[27]The blue line in Panel (a) shows that this sharp increase in Wisconsin also appears among SNAP recipients, the focus of our analysis, with a 26% increase in certifications for the SNAP target group.

[28]This comes from state-level quarterly data from the US Department of Labor rather than WOTC micro-data as in Panel (a). Within this aggregated data, we are only able to look at applications rather than certifications.

increased markedly and persistently in Wisconsin following its adoption of eWOTC. We examine whether this increase in certifications also caused increases in employment and earnings of eligible SNAP recipients relative to ineligible.

We adopt the same event study specification as in Equation (2.1), but around June 2017 rather than January 2007. We again define our treatment and control group as individuals receiving SNAP in three of the past five months who were slightly younger and slightly older than 40 (i.e., aged 37.75-39.75 and 42-43, respectively). This specification estimates a fundamentally different treatment effect in that it picks up the effect of participation costs (as opposed to the effect of the existence of the program). The policy may induce responses among a different set of firms, and so have a potentially different effect on employment. This is particularly relevant for policy since one of the most commonly cited reasons for low uptake of WOTC is the cost of filing (Hunt et al., 2018); we can measure whether the natural policy response to this problem – lowering those costs – has desirable effects. Our results also serve as a policy evaluation of the consequences of switching to electronic filing, which is relevant to other states considering adoption.

Figure 2.6 and Table 2.4 analyzes the effect of electronic filing on the same set of outcomes as previously. As in the other two analyses, we find no evidence that WOTC eligibility affects labor market outcomes (columns (1) to (3)), public assistance usage (column (5)), or criminal activity (column (6)), despite an effect on WOTC certifications.[29] This provides further evidence against the efficacy of wage subsides for promoting employment among this population, as well as suggesting that one policy idea for improving WOTC – lowering access costs – is unlikely to achieve desired outcomes; instead, it just makes it easier for firms to collect the subsidies.

2.5.4 Combining the estimates

Although the individual estimates are typically quite precise, combining the estimates provides even more precision. Since the samples generating the estimates are not necessarily

[29] The coefficient on earnings is negative and just statistically significant at the 10% level. Given the lack of effect on other labor market related variables and that the sign of the coefficient is in the opposite direction from what one would expect, our interpretation is that this is likely to be statistical noise that will naturally occur when running a large number of regressions.

independent (e.g., the same individual could appear in both the RD and SNAP expansion DID samples), we construct confidence intervals around combined estimates via block bootstrapping at the individual level.[30] The combined point estimates are 0.0011 for employment (se=0.00171) and 0.00066 (se=0.00084) for new hires. As a benchmark of precision, these null estimates on employment are an order of magnitude more precise than two recent papers in top economics journals whose main finding is a null effect on employment (Cengiz et al., 2019; Garin et al., 2025).

Another useful estimate is the extent of inframarginality, i.e., among the workers for whom the firm receives a subsidy, what fraction would they have hired even in the absence of the subsidy? We estimate the fraction of subsidized hires whose hiring was caused by WOTC by dividing the estimate for being newly hired into a job by the estimated effect on being hired into a WOTC-subsidized job. The denominator measures how much the variation we exploit increases an individual's likelihood of being hired into a WOTC subsidized job, while the numerator is our estimate of how much hiring likelihood increases. For example, if the likelihood of being hired into a WOTC job increased by 1%, and the likelihood of being hired at all increased by 0.5%, this would imply that half of all WOTC hires were hired as a result of WOTC.[31] Another way of thinking about this estimate is as a normalization – each of our three empirical approaches induces a different increase in WOTC certifications, and so dividing by this generates comparable estimates across strategies.

For each of our three main empirical strategies, we take the estimated effect on the individually getting hired into a job and then divide by the estimated effect on being newly hired into a WOTC job.[32] The precision weighted average of these three estimated ratios is 0.0331. This estimate would imply that 96.7% of hires are inframarginal, and we cannot

[30]In practice, we draw a random sample (with replacement) of Wisconsin residents and then re-run each of the three empirical approaches within this sample. We then construct precision-weighted averages of the estimates, since such weighting approximates the weighting if the estimates were combined using seemingly unrelated regression (Wooldridge, 2010).

[31]While the ratio is likely to be between zero and one, it could in theory even be greater than one. For example, suppose that a firm applies for and receives WOTC for some of its hires but does not screen on WOTC when making hiring decisions. It may respond to an increase in the pool of WOTC-eligible workers (e.g., due to the 2007 SNAP expansion) by beginning to favor WOTC-eligible job applicants at the time of hiring, which could increase the hiring of WOTC eligible workers by more than the marginal increase in certifications the the firm receives due to expansion.

[32]We focus on the outcome of hiring rather than employment since using effects on employment will implicitly weight each hiring episode by length of employment.

reject the null hypothesis that all WOTC hires are inframarginal.[33] This analysis further reinforces that WOTC has little to no effect on employment of WOTC-eligible workers.

2.6 Firm-level analysis

The previous analyses were conducted at the level of the worker rather than the firm. This section re-analyzes the data at the firm-level to examine dynamic and aggregate effects that cannot be tested in the worker-level analysis. For example, it may take time for firms to adjust hiring practices after initially adopting WOTC, which could mute the estimated treatment effects when only looking at the worker side. Alternatively, it could be that there are effects on hiring immediately after adoption, when WOTC is most salient, but these fade away over time. Either dynamic would suggest policy interventions that could improve the efficacy of WOTC: for example, the latter scenario would indicate that increasing salience of WOTC to hiring managers is the key to program improvement.

2.6.1 Firm-level analysis on employment outcomes

Our main firm-level empirical approach leverages the staggered timing over which firms first adopt WOTC. We use an event-study framework, centering our estimates around the date at which firms submitted their first WOTC applications.[34] This allows us to measure how firms' hiring practices evolve as they begin to receive WOTC subsidies for eligible workers. Our focus is on how adopters change their hiring of eligible workers, i.e., those receiving SNAP benefits and slightly below the age eligibility cutoff age of 40, as compared to just ineligible workers, i.e., those receiving SNAP benefits and just above the age of 40, relative to firms that had not yet applied for WOTC. This approach is effectively a triple-difference approach, exploiting differences across firm, time, and type of worker. Examining changes in the *relative* hiring of these highly substitutable sets of workers isolates variation caused by adoption of WOTC rather than other changes that adopting firms may be experiencing relative to non-

[33] As seen in Appendix Figure B11, most of the mass of the distribution is close to zero or even negative.

[34] The WOTC micro-data begins in 2005, meaning that we cannot observe whether a firm submitted a WOTC applications prior to that. To deal with this issue, we restrict adopters to be the set of firms whose first WOTC application is filed in 2007 or later, as the lack of applications between 2005 and 2007 suggests that this was likely their first application.

adopting firms around the time of expansion (e.g., firms may tend to adopt WOTC during an expansion in hiring). The identifying assumption is that firms adopting WOTC do not simultaneously experience shocks to their demand for labor for SNAP recipients just below age 40 as compared to those just above age 40. We provide evidence in Section 2.7 that firms view these workers as highly substitutable, so this appears to be a reasonable assumption.[35]

To analyze the dynamic effects of staggered WOTC adoption on outcomes $y_{f,t}$, we implement an event-study specification following Callaway and Sant'Anna (2021). Specifically, we estimate the average treatment effect on the treated (ATT) for firms at each year relative to adoption.[36] The estimating equation is given by:

$$y_{f,t} = \alpha_f + \lambda_t + \sum_{\ell \neq -1} \delta_\ell D^\ell_{f,t} + \varepsilon_{f,t} \tag{2.4}$$

Here, α_f and λ_t denote firm and year fixed effects, respectively. The variable $D^\ell_{f,t}$ is an indicator that equals one if firm f adopted the policy exactly ℓ periods ago by year t, and zero otherwise. The coefficients δ_ℓ represent the dynamic average treatment effects at each period relative to adoption, capturing the temporal evolution of the policy's impact. Standard errors are clustered at the firm level. We then estimate an overall effect that aggregates across the five years before and after adoption.

In Panel (a) of Figure 2.7, we define the outcome variable $y_{f,t}$ as the number of hires by firm f in year t who are receiving SNAP and are aged 38 or 39 at the time of hiring (i.e., eligible for WOTC) minus the number of hires who receive SNAP and are aged 40 or 41 (i.e., just ineligible for WOTC). In Panel (c), we instead use SNAP recipients aged 35 to 39 as the eligible group and those aged 40 to 44 as the ineligible group to test robustness to a wider age band and improve statistical power. In both cases, we find no increase in the relative demand for WOTC eligible workers despite the economic incentive to shift hiring towards subsidized workers. Table 2.5 provides estimates of the aggregated ATT: the point estimates are small and precisely estimated relative to the average number of hires in a period. These results

[35]The most likely issue would be if firms decide to adopt WOTC because of an unrelated positive shock to their labor demand for SNAP recipients below age 40. That would tend to bias our estimates upwards (i.e., towards finding a positive effect on hiring), but we estimate precise null effects on hiring so this does not appear to be a major problem.

[36]This analysis aggregates the data to the year-level so that the method of Callaway and Sant'Anna (2021) can be implemented, as quarterly data creates too many comparisons comparison groups.

reinforce findings from the analyses in Section 2.5 that firms have limited responsiveness to targeted hiring subsidies.

One concern with this approach is the extent to which WOTC adoption generates a meaningful difference in labor costs when hiring eligible as compared to ineligible workers: for example, adopting firms may quickly discontinue WOTC usage or rarely file WOTC applications. To test this, we redefine the dependent variable $y_{f,t}$ to be equal to the difference in the number of WOTC certifications for just eligible hires and certifications for just ineligible workers. Using the same narrow and wide age bands, Panels (b) and (d) of Figure 2.7 instead show that firms experience a sharp and sustained increase in WOTC usage following adoption, meaning that the subsidy indeed lowers the effective labor cost for eligible hires relative to ineligible.

Another question is whether WOTC may induce aggregate effects at the level of the firm that an individual-level analysis would miss. For example, firms may face pay equity or other internal wage setting constraints that prevent them from raising the wages of individual workers, and so they respond by raising the wages of all entry-level hires (Saez et al., 2019). They may also expand the total number of workers in the firm.

To test for such effects, we exploit variation in the extent to which firms were exposed to the 2007 WOTC expansion. For each industry (at the 4-digit NAICS code level), we calculate the fraction of the industry's hires in the pre-2007 period who received SNAP and were between the ages of 25-39, i.e., who would become eligible for WOTC after the expansion. We then use a difference-in-differences design to test whether firms who have adopted WOTC and are in an industry in the top tercile of exposure (i.e., had a large expected positive shock from the expansion) increase their payroll (potentially due to offering higher wages) or workforce as compared to firms who have adopted WOTC but are in the bottom tercile of exposure.[37] Column (1) of Appendix Table B13 and Panel A of Appendix Table **??** indeed find that more heavily exposed firms have a greater influx of WOTC hires, averaging over 3 additional WOTC hires per quarter. This is similar using a measure of WOTC hires per total hires (column 2 and Panel B). However, the remaining columns of the

[37]We reweight the payroll-by-hires distribution in the bottom tercile group to match that of the top tercile group for better comparability.

table and panels of the figure find little evidence that this translates into additional total hiring (column 3), total payroll (column 4) or starting hourly wages for workers (column 5).[38] Note that even though WOTC-adopting firms file applications for a large fraction of their hires (11%), the subsidy amount is still only a small fraction of its wage bill because WOTC workers earn such low wages. Among WOTC using firms, the WOTC subsidy is equal to 0.17% of their total wage bill, equivalent to roughly a third of the state unemployment tax rate. Given this, it is sensible that any firm-level effects would be muted.

2.6.2 Do firms churn through WOTC-eligible workers?

The short-term nature of WOTC-subsidized employment raises concerns that WOTC may give firms an incentive to "churn" workers, i.e., fire subsidized employees after the subsidy is exhausted to hire a new batch of subsidized workers (Corwin, 2022). As a result, the existence of WOTC could hinder workers from finding stable jobs, including workers who are ineligible for the subsidies. This could be true even if firms do not account for WOTC in hiring decisions, as the results of the previous sections would indicate, where WOTC status may be more salient after hiring and certification are complete.

We test for churning behavior using hours- and earnings-based discontinuities built into the program. The WOTC subsidy is exhausted at \$2400 per worker, which typically corresponds to \$6000 in worker earnings.[39] Churning firms would be expected to terminate workers after reaching that threshold, generating bunching in earnings just above \$6000 for WOTC-subsidized employment stints. They might also not bunch, but exhibit a change in the separation rate at around this threshold: for example, if terminations are often the result of worker infractions (e.g., tardiness), the threshold for termination may drop once the subsidy is exhausted.

To test for these behaviors, we measure earnings and hours worked in WOTC subsidized job stints. Panel (a) of Figure 2.8 plots the distribution of total earnings prior to separation for workers whose employment was subsidized by WOTC (and were eligible for

[38]For this last column, we take the average starting hourly wage for WOTC-hired workers from the WOTC data.

[39]At the 40% subsidy rate, the subsidy would be exhausted at \$6000 in earnings. If a worker earned an hourly rate of above \$24 per hour, then they would exhaust the subsidy at the 25% rate and a higher total earnings amount, but only 0.2% of WOTC subsidized workers have hourly rates that high in our data.

WOTC through SNAP). There is no obvious discontinuity in either total separations or separation rate at \$6000. However, it could be that churning behavior is concentrated among certain types of firms that are most attentive to WOTC. Appendix Figure B12 examines the distribution of worker earnings prior to separation for WOTC-subsidized workers who were employed at the top 50 WOTC-using firms (panel a), firms with high success rates in applying for WOTC (panel b), and employment services firms (panel c), but again finds no evidence of firms responding to subsidy exhaustion. The (lack of) relationship is also present when we split the sample by year of employment (panels d and e), where there is no evidence of churning within more recent years when firms may be more attentive to WOTC or have better payroll processing technology that makes churning easier. While there is no obvious change around \$6000 visually, formal tests also find no evidence of either bunching above the threshold or changes in the rate of separation around this threshold.[40]

The design of the program also provides two additional places to test for churning. As seen in Appendix Figure B1, the subsidy amount increases discontinuously after workers have been in the job for 120 hours (subsidy increases from 0% to 25%) or 400 hours (subsidy increases from 25% to 40%): for a worker earning the median wage of \$9 per hour, firms would get \$540 more in subsidies for keeping them for 400 hours as compared to 399. Churning firms may therefore exhibit bunching to the right of those thresholds, where they keep workers for slightly longer in order to gain the subsidy. Panel (b) of Figure 2.8 plots the distribution of total hours worked in a WOTC subsidized job, with dashed lines marking 120 and 400 hours. The distribution is smooth over those thresholds, with no obvious bunching.[41] Appendix Figure B13 also finds this to be the case when focusing on particular firms that may be more attuned to WOTC or over more recent years.

Based on these analyses, we conclude that firms do not exhibit churning behavior in

[40]For the first test, we test for a change in the density of earnings at \$6000 with a simple local polynomial density estimator from (Cattaneo et al., 2020) ($p = 0.86$). To test for a change in rate of separation, Appendix Table B11 estimates a hazard model comparing the hazard rate of separation for WOTC hires as compared to hires at WOTC-using firms for whom a WOTC application was not filed. We find no differential change in the hazard rate of separation around \$6000 in earnings for WOTC hires relative to non-WOTC hires.

[41]This figure is similar to Figure 4 in Hamersma (2011), which plots estimated hours worked in WOTC-subsidized jobs between 1998 and 2003 to visually examine if there is bunching around 120 and 400 hours (although not for changes around subsidy exhaustion). We augment her analysis by including data from after 2003, conducting formal statistical tests as is permitted by our much larger sample size (over 430,000 certified work stints as compared to 8,908), and examining heterogeneity across types of firms.

response to WOTC. It could be that managers are not aware of the benefit schedule for the firm, and so it would be difficult for them to target these discontinuities. Another possibility is that recruitment or training costs are sufficiently high that firms prefer to keep trained workers rather than churn through them for increased WOTC wage subsidies. That seems particularly plausible given that firms do not appear to target hiring towards WOTC-eligible workers, meaning that the odds of getting WOTC subsidies for a replacement worker may not even be that high.

2.6.3 Firm characteristics

Our analysis indicates that the benefits of WOTC accrue to employers rather than workers. This naturally raises questions about what types of firms benefit, and whether it is desirable for policymakers to subsidize their operations.[42] We begin by examining the distribution of benefits among firms. We rank all firms in the data in order of the number of WOTC certifications they receive between 2005 and 2019 in Wisconsin, with rank 1 being the firm with the most certifications, and so on. In Figure 2.9, the blue line plots the cumulative fraction of WOTC certifications accruing to firms of rank between 1 and N for the top 300 firms. Benefits accrue to a small number of firms, where over half of all certifications went to just 48 firms. However, it could be that this is not a disproportionate fraction of rents, but that these firms also employ a large fraction of all workers. The red line instead plots the cumulative fraction of hires for these firms, finding that they receive 4 to 5 times the amount of tax credit expected given their share of employment: for example, the top 48 firms accounted for just 9% of all hires over the period.

Given this high concentration of WOTC certifications among firms, we focus on understanding the 50 firms with the most WOTC certifications over this period. These firms are large, where we observe an average quarterly payroll of over $20 million in the years in which they receive WOTC certifications. For additional details, we use data from FOIA requests made to the Wisconsin Department of Workforce Development, which provide the total number of WOTC applications and certifications by firm between 2018 and 2020 (Corwin,

[42]Benefits could also accrue to consumers if the subsidy is passed through to prices. We are unable to measure this with our data, but it is unclear whether this would be expected given price rigidities and the typically short-run nature of WOTC employment.

2022).[43] Just over half of the top 50 firms (52%) are temporary staffing services agencies.[44] Around a quarter (24%) are large publicly traded firms, where these firms have a median market capitalization at over $30 billion, with only one having a market capitalization below $1 billion. Most of the remaining privately held firms are large franchisee operators of national fast food chains (16%) or large private companies with annual revenues in excess of $5 billion annually (4%). The high number of temporary staffing firms is notable given experimental work showing that employment in temporary-help positions may actually decrease future employment and earnings (Autor and Houseman, 2010). As a result, subsidizing those type of firms is unlikely to achieve long-run employment gains for disadvantaged workers.

2.7 Mechanisms

These findings present a puzzle: why do profit maximizing firms fail to respond to subsidies like WOTC? This is particularly puzzling given that these firms *do* apply to receive the wage subsidies after hiring eligible workers, so lack of knowledge about the program cannot explain their lack of response. Determining the mechanisms underlying these results is essential in considering the extent to which wage subsidies could benefit targeted workers in other contexts. This section evaluates three possible mechanisms.

Standard model of labor supply and demand. We first consider whether these results can be reconciled within a standard model of labor supply and demand. In this model, the wage subsidy shifts up the firm's labor demand curve for subsidized workers, as firms now have a lower effective cost of their services. The lack of change in both employment and wages could only be rationalized by perfectly inelastic demand for these workers: firm labor demand is fixed, and the full value of the subsidy accrues to firms as the more inelastic side of the market.[45] Perfect inelasticity in demand for these workers implies a lack of substitutability

[43]The Wisconsin Administrative Data Core data is anonymized and so researchers using it do not have access to the identity of particular firms. We are grateful to Emily Corwin for sharing this data with us (see Corwin (2022) for details).

[44]Because workers at temporary staffing agencies are directly employed by the staffing agency (and not the firm to which they are sent), it is the staffing agency that receives WOTC subsidies. However, there could be some pass-through to the contract with the receiving firm.

[45]If, instead, labor supply were perfectly inelastic and labor demand was downwards sloping in wages, then the subsidy would fully pass through to the worker as an increase in wages. If labor supply were perfectly

between eligible and ineligible workers. That would be surprising given these workers are often quite similar: for example, our empirical approaches typically compare labor outcomes of SNAP recipients who are just above and below a particular age cutoff (such as age 40). For this model to be correct, firms must not be able or willing to substitute workers just above age 40 (ineligible) with those just below (eligible), which seems unlikely.

To empirically test the substitutability of workers who are just-eligible and just-ineligible for WOTC, we examine replacement hiring after exits of WOTC-eligible and WOTC-ineligible workers. If production processes are such that WOTC-eligible workers are not substitutable with other types of workers (producing perfect inelasticity in demand for their services), the exit of this type of worker from a firm would necessitate the hiring of another of this type of worker as a replacement, as opposed to hiring other types of workers. For each firm, we measure the number of workers who exit the firm in a given quarter t who received SNAP and were between the ages of 37-39 at the time of hiring (eligible exits), as well as those who received SNAP and were between the ages of 42-44 (ineligible exits). Columns (1) and (2) of Appendix Table B12 report how the number of new hires of eligible workers (SNAP recipients aged 37-39) in quarter t responds to these exit variables, while columns (3) and (4) does the same with new hires of ineligible workers (SNAP recipients aged 42-44). We include fixed effects for quarter, firm, and total number of exits for the firm in a quarter to account for any firm-specific patterns in separation from and hiring of workers. If the two groups of workers are highly substitutable, we would expect similar coefficients in both specifications, i.e., that the exit of a 37-39 SNAP recipient can be compensated by hiring either a 37-39 or 42-44 year old SNAP recipient (or other types of workers). Unsurprisingly, the coefficients on "Eligible Exits" are quite similar across specifications, consistent with a high degree of substitutability between these types of workers such that perfectly inelastic demand does not explain our results.

Negative signal value of subsidies. Another possible model is that firms infer eligibility for wage subsidies as a negative signal of productivity, as seen in the seminal study of

elastic and labor demand were downwards sloping, the subsidy would induce an increase in employment. If labor demand were perfectly elastic or in intermediate cases, the wage subsidy should lead to an increase in both wages and employment.

Burtless (1985). If this negative signal exactly offsets the value of the subsidy to the firm (in expectation), then this could generate a lack of firm response to WOTC eligibility.

Our results point against this hypothesis. In particular, suppose that the negative signal exactly offset the net value of the subsidy prior to the 2017 introduction of e-filing. The reduction in (implicit) application costs starting in 2017 should have increased the hiring of WOTC-eligible workers among firms for which the expected benefit of the subsidy and the negative signal value of WOTC eligibility were previously equal. However, our analysis in Section 2.5.3 does not find any employment response. The negative signaling mechanism thus does not appear to explain our results. The next section will also provide evidence against this explanation, showing that those making hiring decisions typically do not know who is eligible for WOTC.

Information about eligibility. Another potential reason why firms may not respond to WOTC would be if hiring decision makers do not observe WOTC eligibility at the time of hiring decisions, such as if there are high screening costs or other constraints on information acquisition. To investigate this explanation and to understand how WOTC integrates into hiring processes more generally, we combine three pieces of evidence: (1) original data on job applications for openings at a sample of WOTC-using firms; (2) an audit study experiment on how firms respond to signals of WOTC eligibility among job applicants; and (3) an online survey of individuals with hiring experience at WOTC-using firms.

We first collect data from a random sample of Wisconsin firms that use WOTC. Taking data from FOIA requests to the Wisconsin Department of Workforce Development, we observe each WOTC application and certification between 2018 and 2020, including the name of the associated firm. We drew a random sample of these firms using probability proportionate to size sampling, searched for current job postings by the firm (available for 83% of firms), and submitted job applications for a position in the firm.[46] We recorded data on characteristics of the job application process, including whether the application asked WOTC-related eligibility questions and, if so, the specificity of the questions.

Table 2.6 reports summary statistics from these applications. One striking fact from

[46]In cases where it was possible to apply for multiple positions, we selected the position that seemed most likely to be relevant for WOTC based on the posted wage, desired skills, or other attributes of the job.

this table is the relatively small number of job applications that ask questions about WOTC eligibility. Firms do not typically have access to government data that could indicate whether a potential hire is eligible for WOTC, and so must ask applicants about criteria related to eligibility (e.g., receipt of SNAP benefits) to identify WOTC-eligible hires.[47,48] However, even among this sample of firms that file for WOTC, only 13% of firms collect information that would allow determination of WOTC eligibility on their job applications. Furthermore, in 44% of these cases, the firm itself does not collect the information; instead, the applicant is taken to the separate website of a hired consultant or the firm's payroll processor, which collects this information. Our discussions with firms indicate that this information is not typically shared with the person making hiring decisions, and as seen in the job application form in panel (a) of Appendix Figure B15, most of the forms collecting WOTC data explicitly note that this will not be shared with the firm. Instead, if the worker is hired, the outside party uses this information to file for WOTC tax credits on behalf of the firm. In many cases, firms that file for WOTC collect information on eligibility after the hiring decision is made, such as during onboarding, and use this information to file for the wage subsidy. Overall, our job application data suggests that the person making hiring decisions for WOTC eligible workers is aware of their eligibility in less than a tenth of cases.

The lack of explicit screening on WOTC eligibility does not necessarily preclude WOTC from having an effect on employment of eligible workers. For example, those responsible for hiring could still respond to observable characteristics related to WOTC eligibility even if they do not know exactly who is eligible.[49] As a test of this possibility, Appendix Table B14 repurposes the audit study data from Kline et al. (2024), which submitted more than 83,000 job applications to the types of entry-level jobs for which WOTC is common. Kline et al. (2024) randomize applicant characteristics such as age, and their replication data provides the name of the firm to which each application was submitted. As with the job application

[47]An exception to this is felony status, which can be determined through a criminal background check. However, individuals with a felony conviction are a relatively small share of all WOTC-eligible individuals (less than 4%). For the most common reason for eligibility – receipt of SNAP benefits – there is not a straightforward way for firms to determine this independently.

[48]Two examples of how firms collect this information can be found in Appendix Figure B15.

[49]Another possible channel of response is that firms using WOTC may advertise more heavily to pools of potential job applicants that are WOTC-eligible. However, the lack of effects on employment in section 2.5 would indicate either that this does not happen, such advertising is ineffective at raising applications, or any increase in applications among the WOTC-eligible is eliminated during screening in the hiring process.

analysis above, we determine which of those firms use WOTC using FOIA data on WOTC applications between 2018 and 2020 that provides the name of the associated firm (Corwin, 2022). Using the same regression discontinuity approach as in Section 2.5.2, column (1) finds that firms that use WOTC are no more likely to call back applicants whose age could make them eligible for WOTC. Column (2) runs the same analysis among non-WOTC using firms to rule out confounds from other possible policies using this age threshold, while column (3) confirms a lack of effect when differencing the estimates in columns (1) and (2) with a differences-in-discontinuities approach. Columns (4) and (5) also find no evidence of effects when running the same specification but restricting to two types of job applicants whose observable characteristics may make firms more likely to think they receive SNAP (Black applicants and applicants without a college degree). Thus firms do not appear to be screening on WOTC in either an explicit or implicit fashion.

These exercises clarify the mechanism underlying the lack of response to WOTC: firms do not set up hiring processes to screen on WOTC, so the person making hiring decisions typically lacks information to respond to WOTC incentives. But this raises an even larger puzzle of why firms do not collect this information and encourage hiring managers to incorporate it in their decisions. To understand this, we conducted surveys with individuals involved in making hiring decisions at firms employing low-wage workers. Respondents were sampled through the Prolific platform and compensated at a rate of $20-30 per hour for taking our survey.[50] Our survey asked questions about hiring practices at their company, and, if relevant, questions about practices around WOTC (e.g., their screening and application procedure) as well as how WOTC eligibility affects their decisionmaking.

We focus on respondents who worked at firms that receive WOTC subsidies but state that their firms do not ask about WOTC prior to the hiring decision. We asked these respondents for the reasons why they do not ask about WOTC, and then to rank these reasons in order of importance. The most common reasons were "fear of discrimination lawsuits" (68% of respondents), "Don't want to scare people off of applying for the position" (38%), and "Hard to explain to applicants why we are asking these questions" (32%).[51] The first option refers to

[50] Only Prolific members who reported working at a private firm and having experience hiring workers that earn less than $20 per hour were selected for the survey.

[51] After these, the next most common reasons were "Don't want to make job applications too long" (21%)

the concern that unsuccessful job applicants may sue the company alleging discriminationary employment practices: for example, the structure of WOTC pushes firms to favor younger applicants (under age 40), which could run afoul of age discrimination legislation. The possibility of having to face such costs as an individual firm generates a strong incentive to not to consider WOTC status at the time of hiring. That is particularly true since introducing explicit screening prior to hiring may only marginally increase the number of WOTC eligible hires relative to simply hiring from the pool of low-wage workers, many of whom will anyway be WOTC eligible.

The second and third most common reasons refer to the possibility that explicit screening for WOTC could make it more difficult to attract applicants. In follow-up phone interviews with hiring managers conducted by the authors, multiple managers noted that WOTC-eligible workers often fear replying in the affirmative to questions about their eligibility as they perceive it could lead to their being screened out. Particularly in high turnover jobs like many of those in which WOTC certifications are common, firms need to continue to attract applicants. It may be that explicit pre-screening on WOTC could deter applications to such an extent that it is preferable to wait until after hiring decisions are made. It may also be preferable to wait since applicants who fear being stigmatized for WOTC eligibility will have less of an incentive to lie about eligibility after the hiring decision has been made.

In countries with active employment discrimination protections, legal liability may present a problem for targeted wage subsidies; even if cases are unlikely, the high cost from legal proceedings could lead to expected costs outweighing expected benefits. Applicant uncertainty about screening questions may also deter potential hires from undergoing the job application process, while high turnover firms need a large pool of applicants. Both of these issues could generate problems for targeted subsidy programs outside of WOTC, but are not prominently discussed in the academic literature. Unfortunately, neither issue has a simple policy remedy. While the Equal Employment Opportunity Commission has issued an opinion that favoring WOTC-eligible applicants does not constitute discriminatory hiring, our consultations with legal experts indicate that this is not binding, particularly in light

and "Too time-consuming for the company to keep track of" (21%).

of a recent Supreme Court decision striking down "Chevron deference".[52] It may be that explicit legislation and/or judicial rulings are needed to resolve the issue, but in the absence of cases, this may not be soon forthcoming. These results suggest that legislators hoping to increase employment of those disadvantaged in the labor market may need to consider means other than targeted wage subsidies.

2.8 Conclusion

Hiring subsidies are a potentially important tool to promote economic mobility for disadvantaged workers, but there is a lack of evidence on their efficacy in the US context. Using data from the Work Opportunity Tax Credit and multiple complementary empirical approaches, we find that these wage subsidies do not promote hiring of target groups. We further find that they are not passed through to worker earnings in any meaningful way, but instead primarily accrue to a small set of large firms.

These findings suggest a need to rethink when and where subsidies can be effective in boosting employment, and in particular, how features of low-wage labor markets affect the efficacy of subsidies. Future work should investigate other means of encouraging employment among these groups, as well as whether hiring subsidy programs like WOTC can be redesigned to achieve this purpose.

[52] This refers to the practice of deferring to the interpretation of relevant federal agencies like the Equal Employment Opportunity Commission in cases of where legislation is unclear, originating in the 1984 Supreme Court case of "Chevron U.S.A., Inc. v. Natural Resources Defense Council, Inc.". In June 2024, the Supreme Court overruled the earlier decision in "Loper Bright Enterprises v. Raimondo".

2.9 Figures and Tables

Figures

Figure 2.1: National Trends in WOTC certification

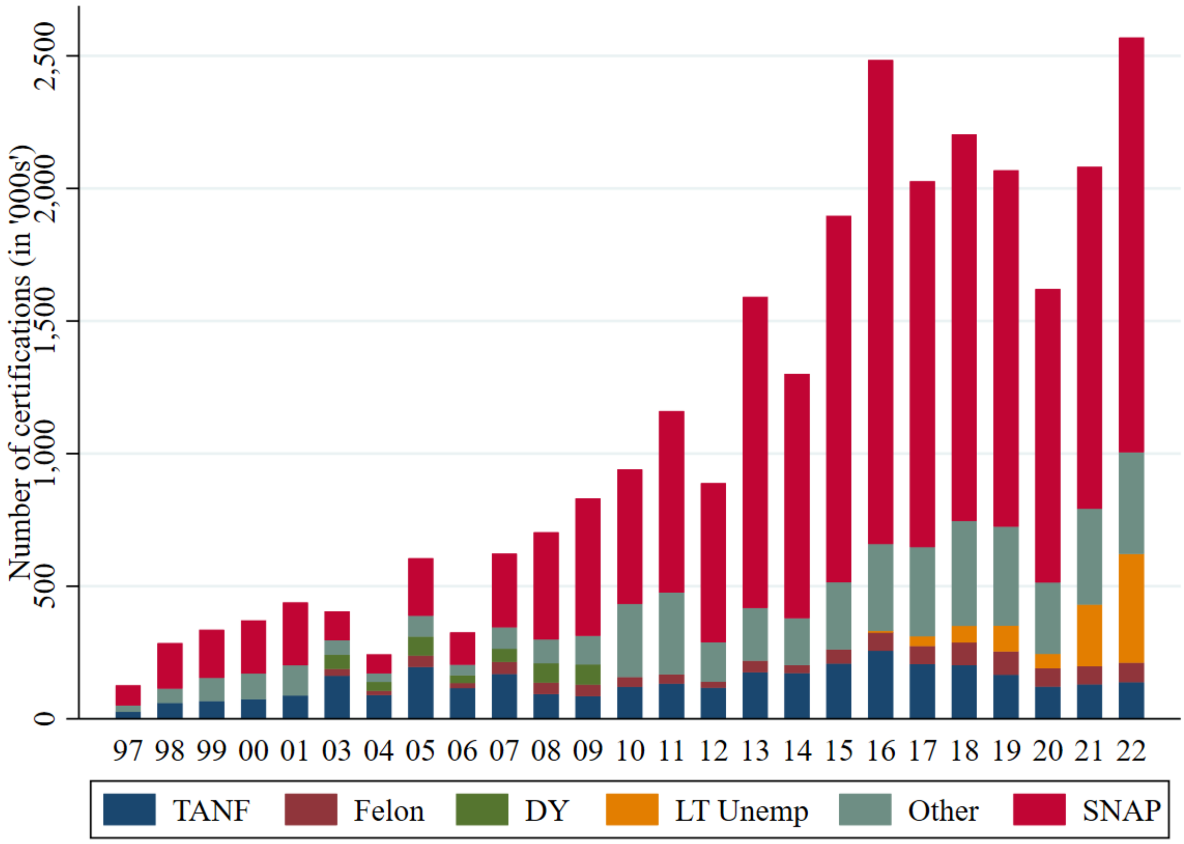

Note: Panel (a) presents the total annual WOTC certifications nationally from 1997-2022 using data from Department of Labor (DoL) annual reports. It splits the certifications by the reason for eligibility, where the Other category combines smaller categories such veterans, youth summer jobs, and long-term TANF. The DoL reports provide number of certifications by year of *certification* rather than *hiring*. There were temporary WOTC lapses in 2004, 2012, and 2014 during which firms could still receive WOTC subsidies for hires made during those periods, but the certifications were not issued until after the lapsation. This inflates the numbers in 2005, 2013, and 2015 due to displacement from the previous years: e.g., the lapsation between January 1, 2004 and October 4, 2004 displaced certifications from 2004 into 2005. The remaining analysis uses individual-level Wisconsin data with date of hiring and so does not have this issue. For 1998, 2000, and 2001, we only observe the national totals, so we use the category percentages from the most recent preceding year with data to estimate the breakdowns. For 1997 to 2002, we do not observe the felon or disconnected youth categories separately, so this is included among "Other".

Figure 2.2: Number of WOTC certifications in Wisconsin (2005-2008)

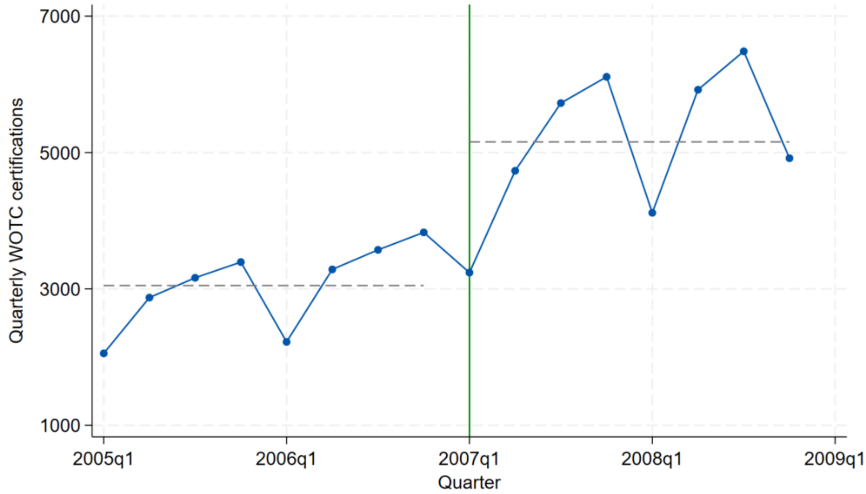

Note: This figure reports total WOTC certifications in Wisconsin by quarter of application from 2005-2009. Gray dashed horizontal lines signify average quarterly certificates over the two-year period.

Figure 2.3: Effect of WOTC on individual outcomes, event study analysis of SNAP expansion

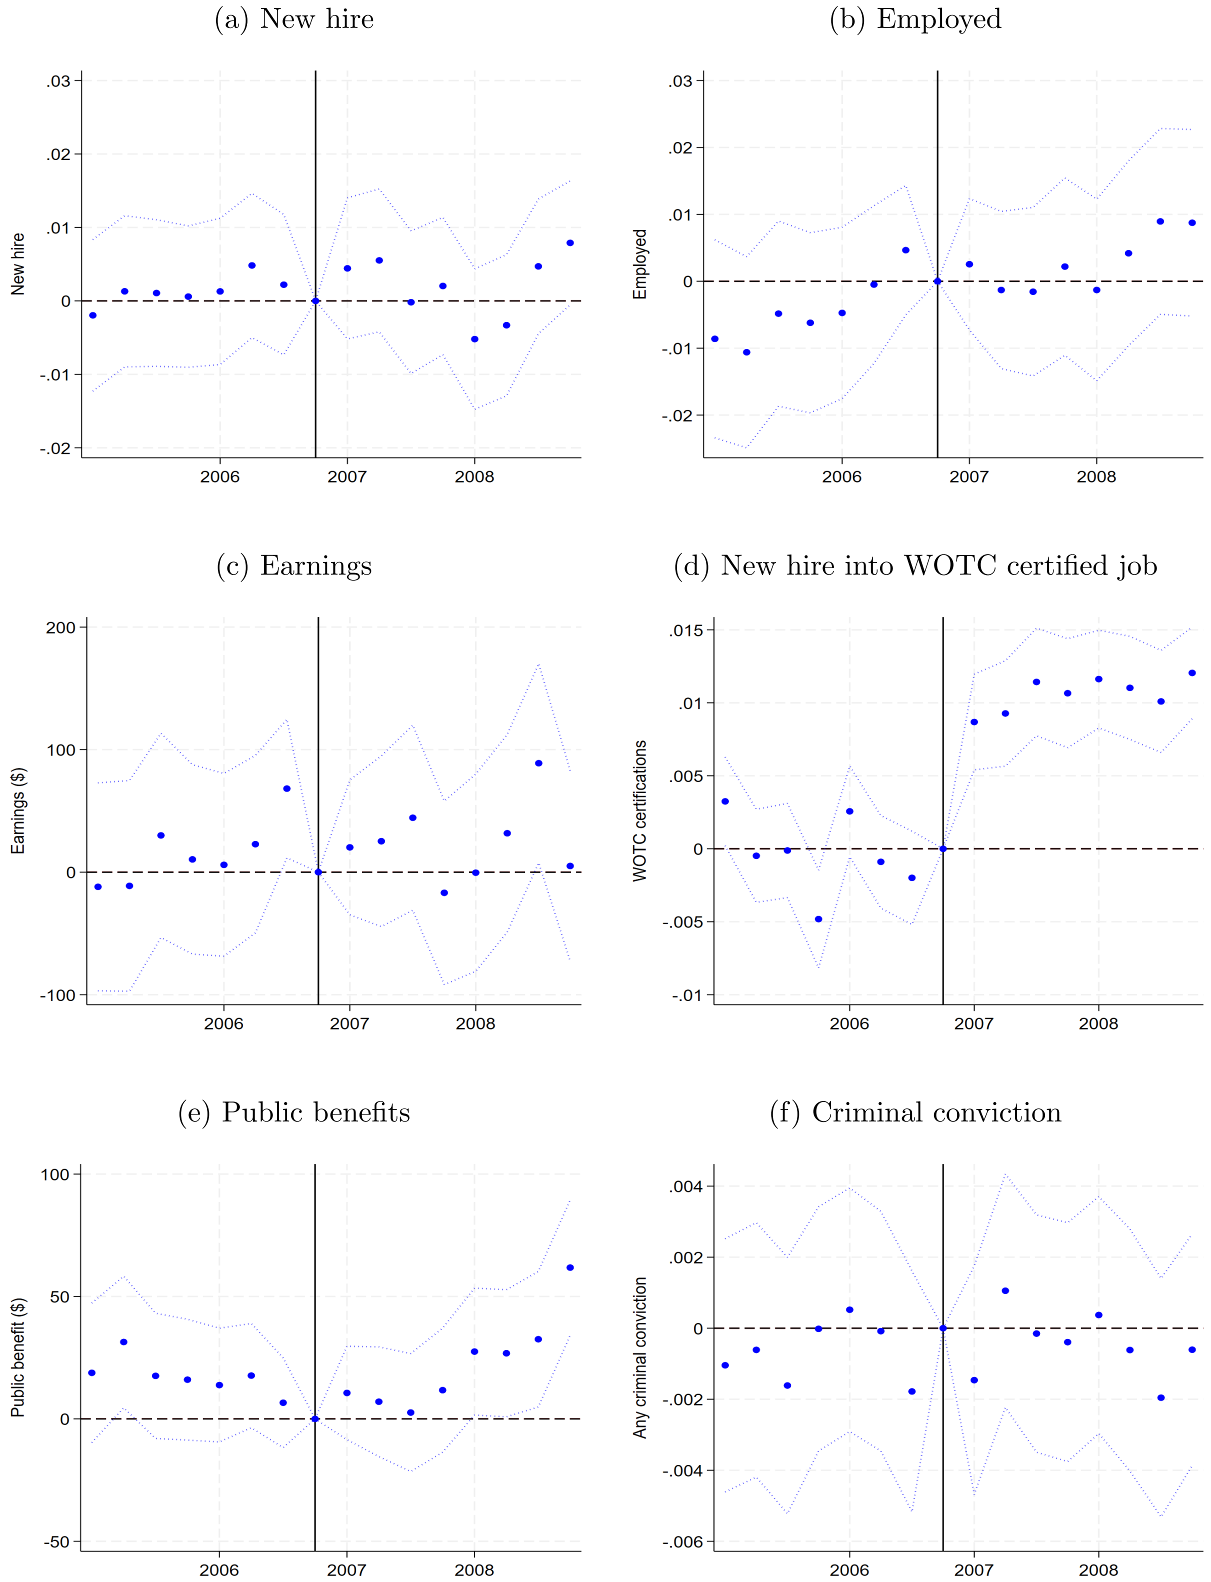

Note: Figure reports stacked event study estimates from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.1) for new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Controls include monthly age fixed effects and quarter-year by age sample fixed effects. Standard errors clustered at the individual level.

Figure 2.4: Effect of WOTC on individual outcomes, RD analysis

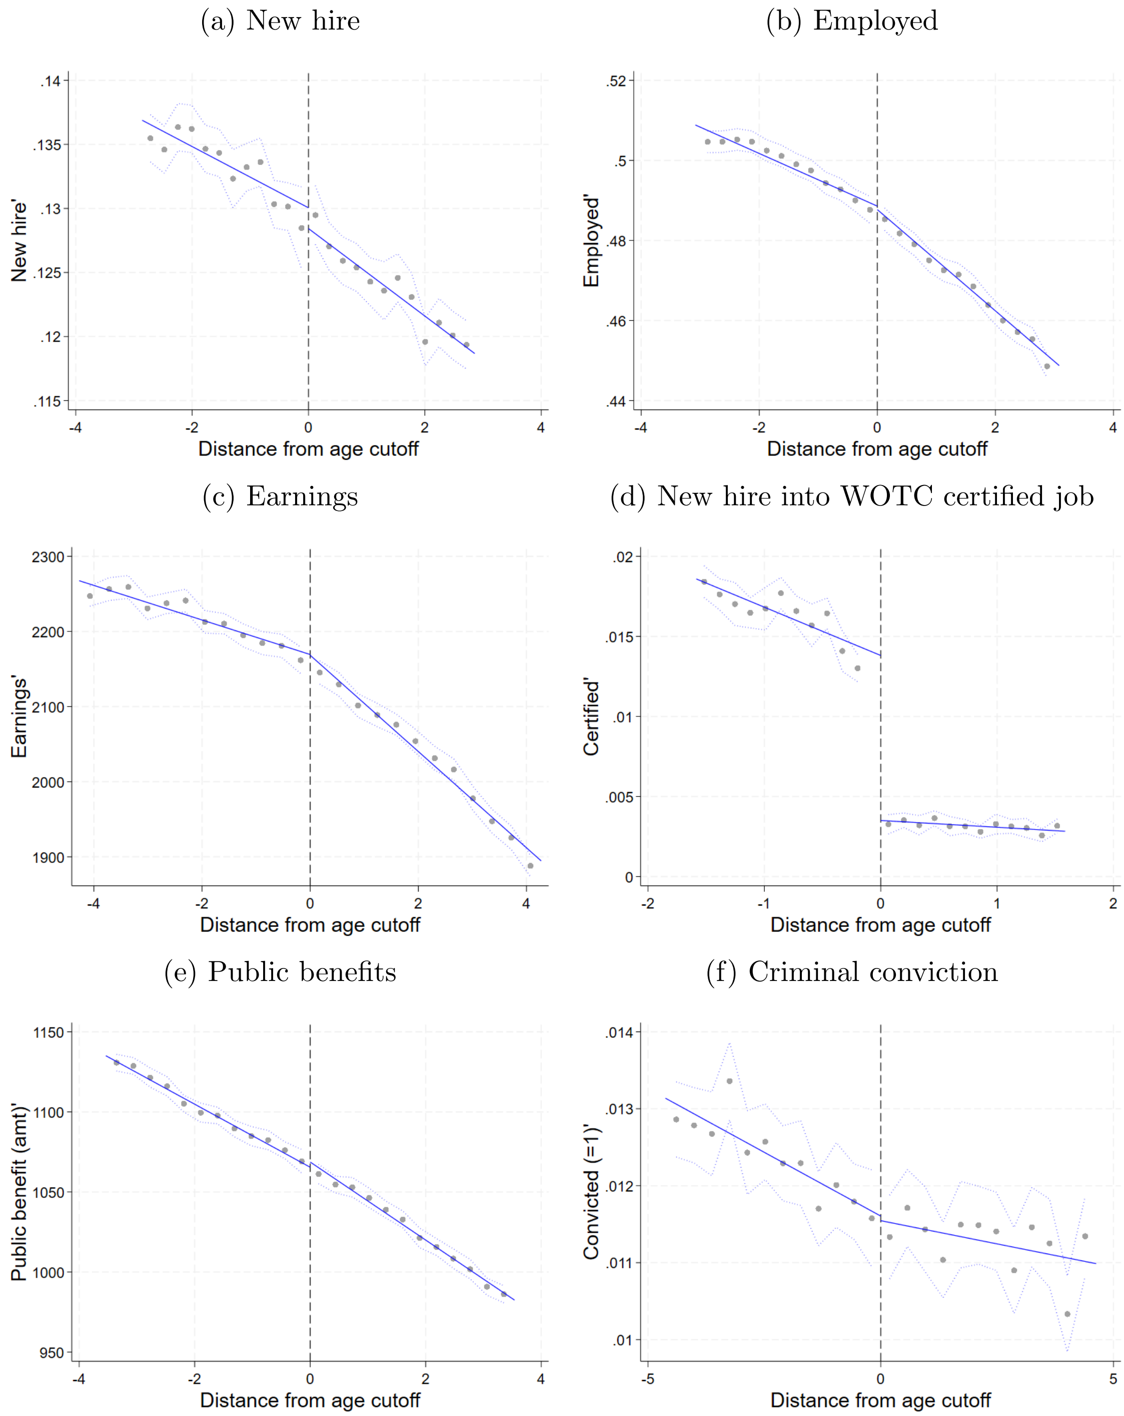

Note: These figures plot the relationship between individual-level outcomes and the difference between the individual's age and the relevant WOTC age eligibility threshold. Individuals whose age is to the left of the threshold are eligible for WOTC, while those to the right of the threshold are not. The bandwidth is based on the optimal bandwidth selection from Cattaneo et al. (2020). Panel (a) examines whether the individual is hired into a new job in that quarter, panel (b) focuses on whether the individual is employed, and panel (c) measures individual earnings in the quarter. Panel (d) plots the first stage relationship between the individual's age and whether they are newly hired in a job with an associated WOTC certification. Panel (e) measures the relationship with the dollar value of three forms of public benefits in a quarter (TANF, Unemployment Insurance, and SNAP), while panel (f) measures if the individual is convicted of a criminal offense in the relevant quarter.

Figure 2.5: WOTC utilization around adoption of electronic filing in Wisconsin

(a) Quarterly WOTC certifications in Wisconsin, 2016-2019

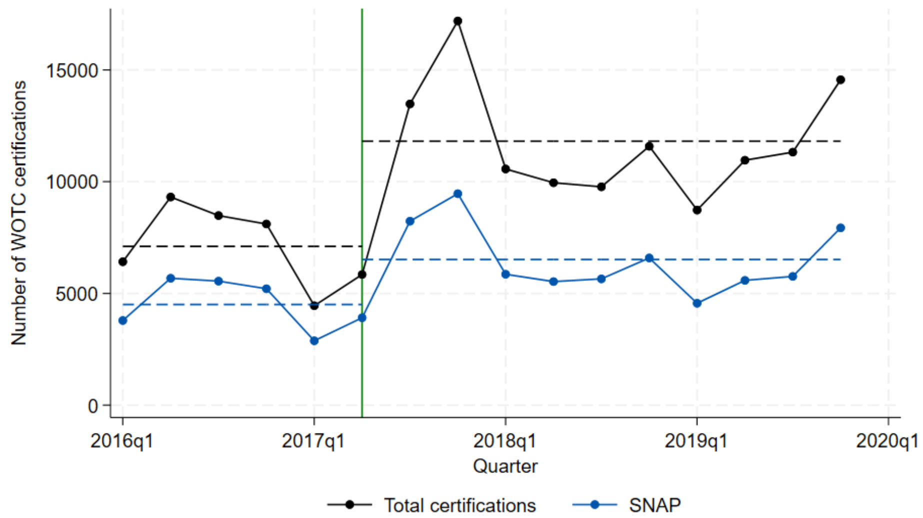

(b) Quarterly WOTC applications in Wisconsin and Minnesota, 2016-2019

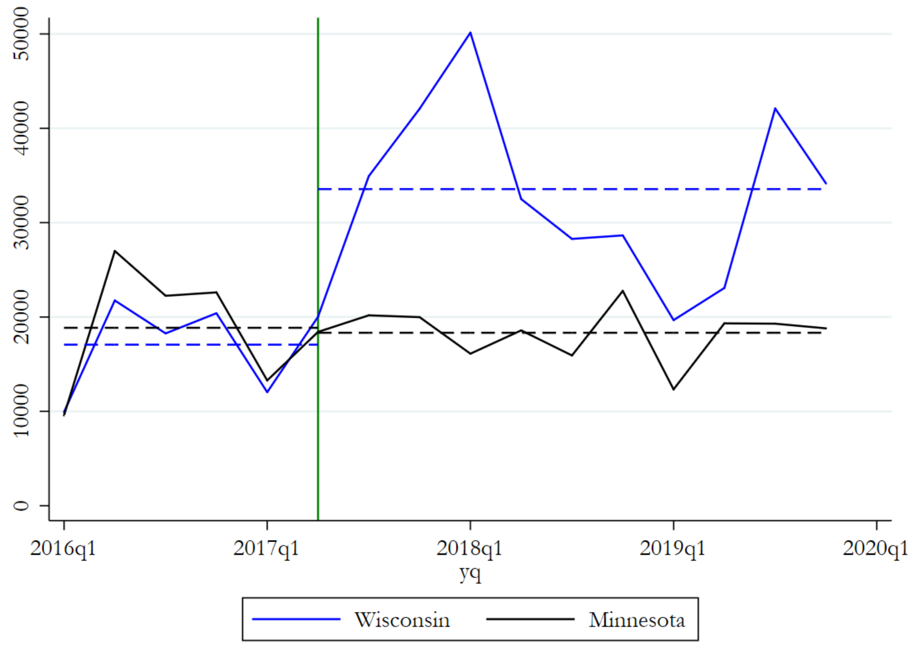

Note: Figure reports total WOTC certifications in Wisconsin by quarter of application from 2015-2019. The vertical green line indicates the date of e-filing adoption. Data is taken from the WOTC micro-data provided to the researchers by the Department of Workforce Development. Panel (b) reports total WOTC applications in Wisconsin and Minnesota by quarter of application from 2016-2019. The vertical green line indicates the date of e-filing adoption in Wisconsin. Data is from the Department of Labor national WOTC data, which is based on quarterly data submitted by state workforce agencies through the the web-based Tax Credit Reporting System. The figures differ due to the different outcome variables (certifications versus applications): for using the national data, it is necessary to use applications as the dependent variable due to the reporting system.

Figure 2.6: Effect of eWOTC on individual outcomes

(a) New hire

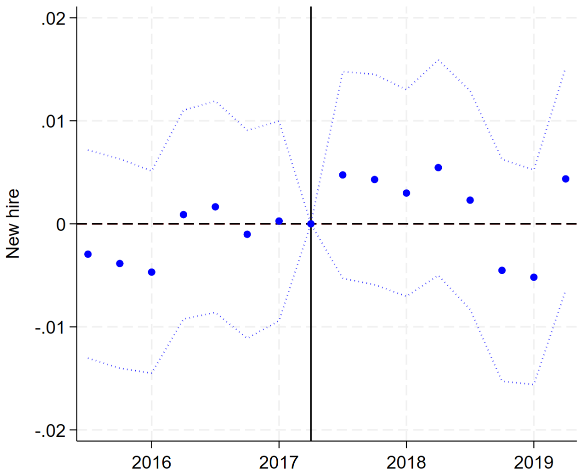

(b) Employed

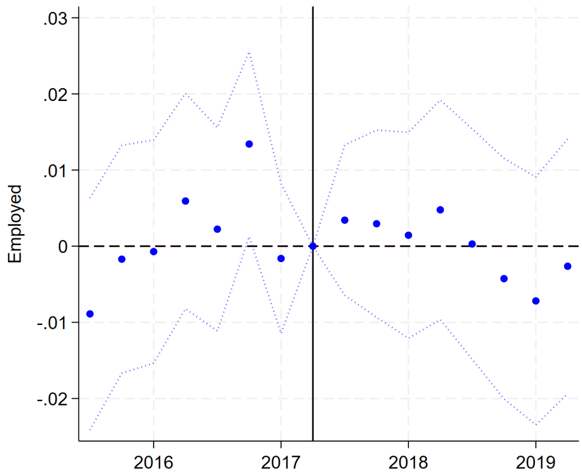

(c) Earnings

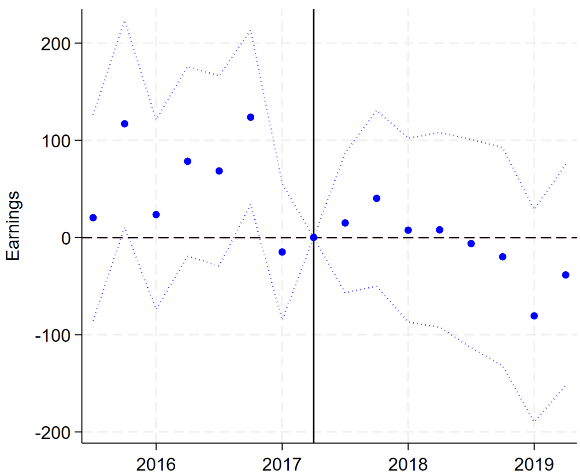

(d) New hire into WOTC certified job

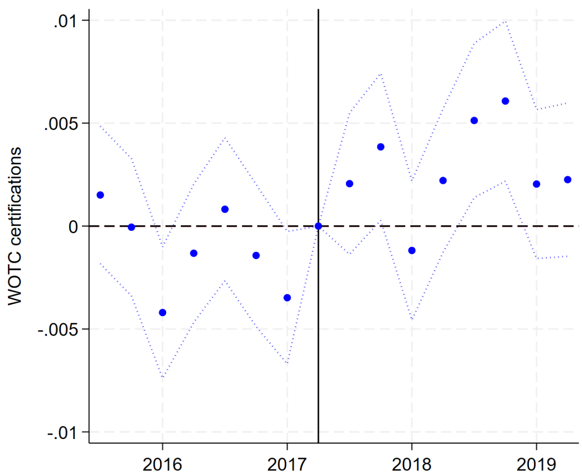

(e) Public benefits

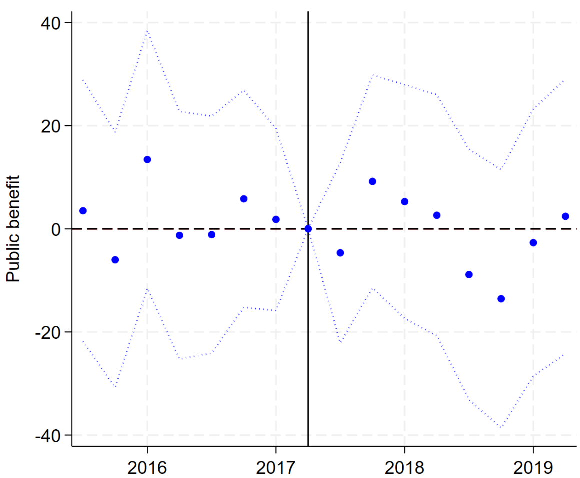

(f) Criminal conviction

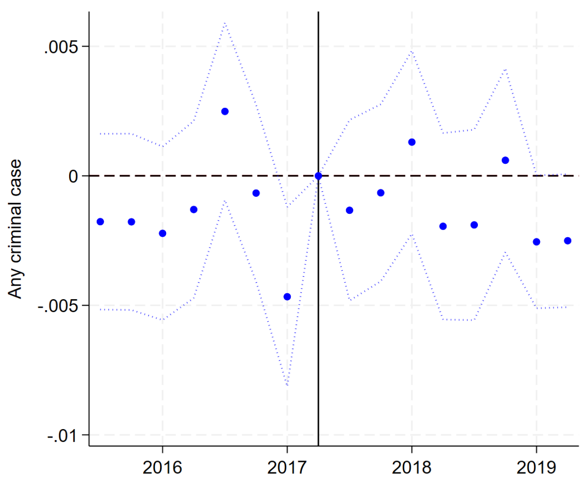

Note: Figure reports event study estimates from the 2017 shift to electronic filing of WOTC applications among our SNAP sample for employment, new hires, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Controls include monthly age fixed effects and quarter-year fixed effects. Standard errors clustered at the individual level.

Figure 2.7: Effect of WOTC on firm-level outcomes

(a) Difference in hires, age 40-41 vs 38-39

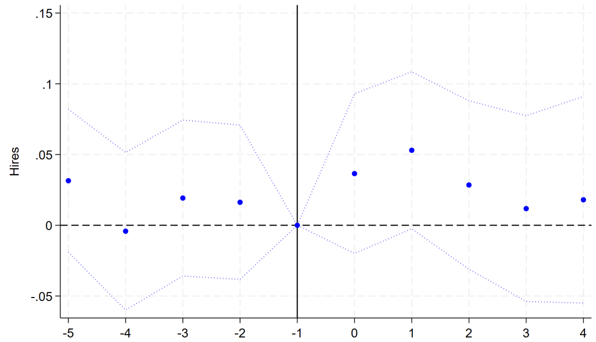

(b) Difference in WOTC certs, age 40-41 vs 38-39

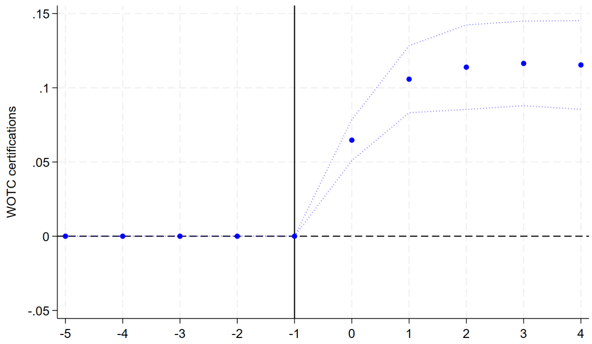

(c) Difference in hires, age 40-44 vs 35-39

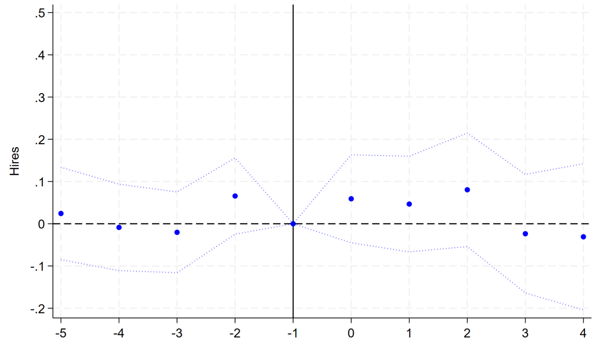

(d) Difference in WOTC certs, age 40-44 vs 35-39

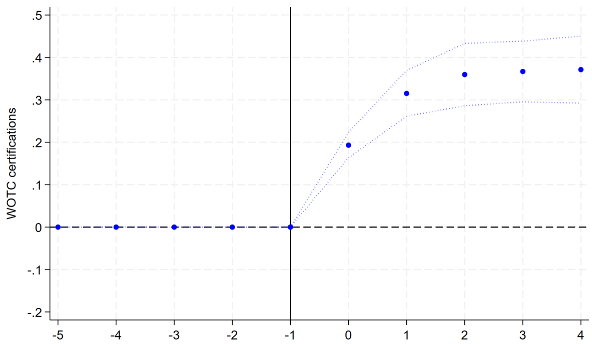

Note: This figure reports the estimated effects of firm-level adoption of WOTC on hiring and utilization of WOTC. The outcomes in the first two panels are the difference in number of hires (WOTC certifications) by the firm of people who are receiving SNAP and are aged 38 or 39 at the time of hiring (i.e., eligible for WOTC) and those who receive SNAP and are aged 40 or 41 (i.e., just ineligible for WOTC). In the third and fourth panels, we instead use SNAP recipients aged 35 to 39 as the eligible group and those aged 40 to 44 as the ineligible group to test robustness to a wider age band. To analyze the dynamic effects of staggered WOTC adoption, we implement an event-study specification following Callaway and Sant'Anna (2021). All regression include firm and year fixed effects. The coefficients estimate the dynamic average treatment effects at each period relative to adoption. Standard errors are clustered at the firm level to account for serial correlation.

Figure 2.8: Histograms of total earnings and hours worked for WOTC-certified workers

(a) Distribution of Earnings

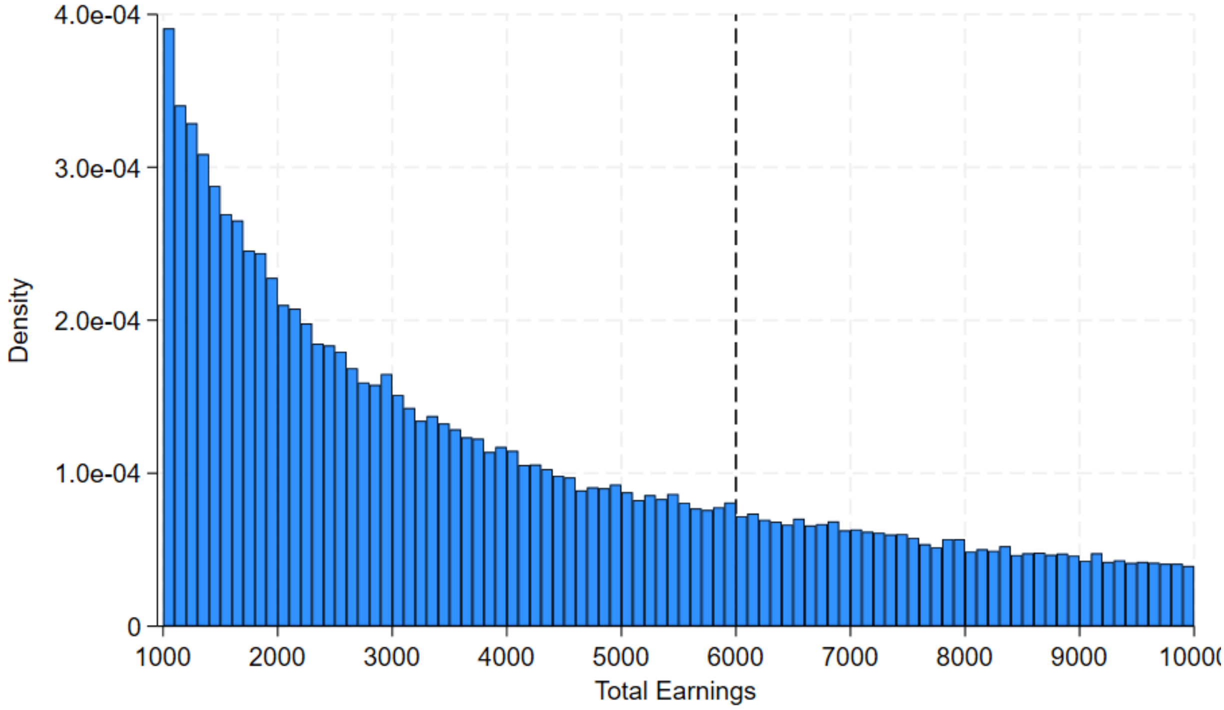

(b) Distribution of Hours Worked

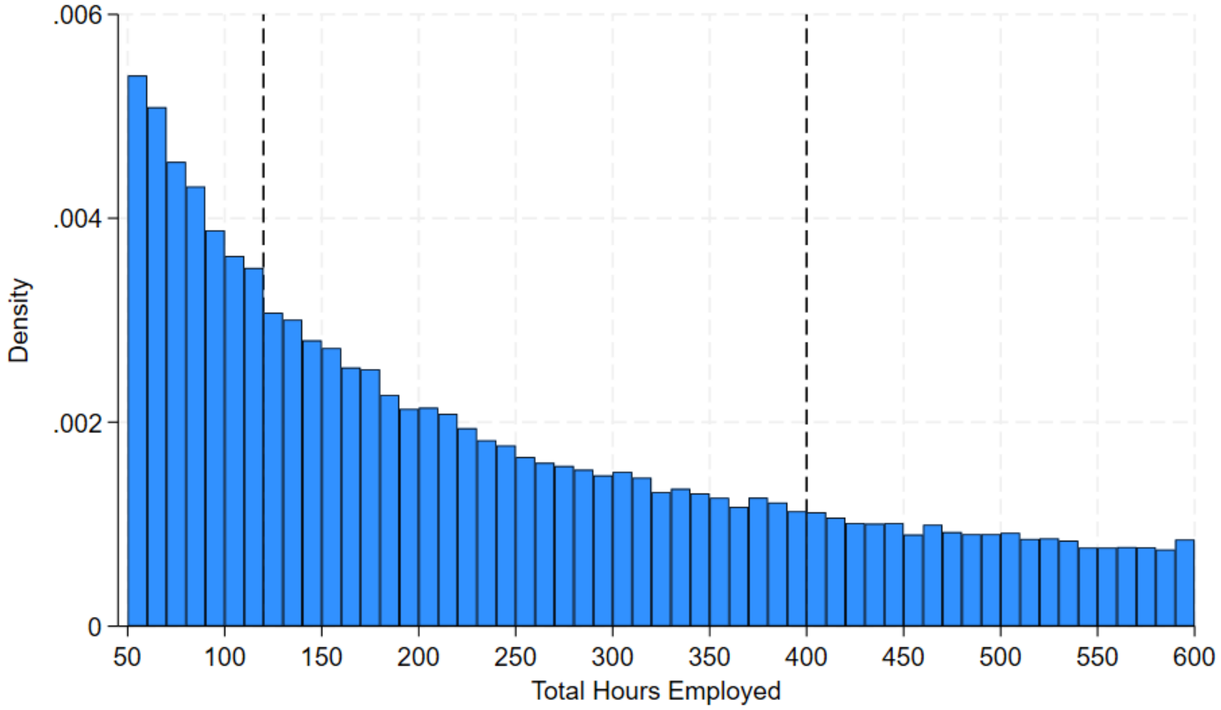

Note: Figures plot the distribution of total earnings and hours worked prior to separation for workers whose hiring was subsidized by WOTC and who were eligible for WOTC through SNAP receipt.

Figure 2.9: Cumulative density function of WOTC certifications by firm

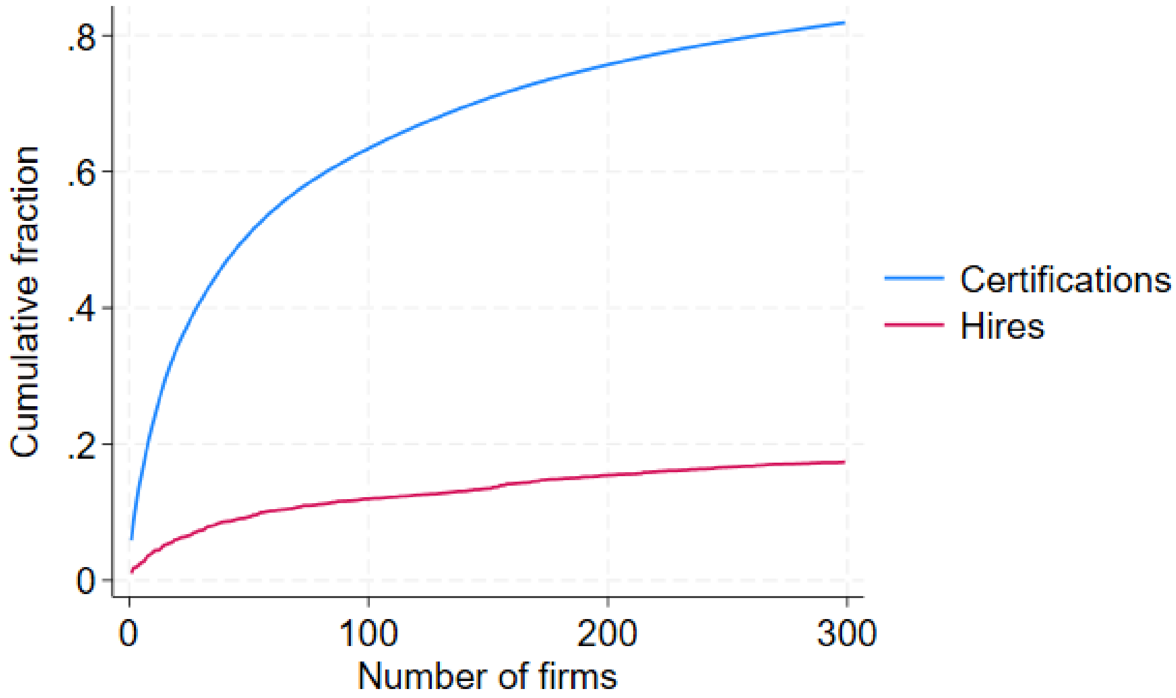

Note: This figure measures the concentration of WOTC certifications among firms. Ranking the firms in terms of WOTC certifications between 2005 and 2019 in Wisconsin, the blue line plots the cumulative fraction of WOTC certifications accruing to firms of rank between 1 and N for the top 300 firms. The red line plots the cumulative fraction of all hires made by those firms over 2005-2019.

Tables

Table 2.1: Summary statistics, individual level

	WOTC applicants	WOTC certified	SNAP sample, 18-39
Characteristics			
Age	29.1	28.2	28.7
Male	0.41	0.36	0.35
White	0.59	0.55	0.63
Black	0.37	0.43	0.32
Hispanic	0.10	0.10	0.11
Born in WI	0.62	0.64	0.64
SNAP beneficiary in quarter	0.63	0.77	0.88
Quarterly SNAP benefits, cond. ($)	945.5	1042.6	1012.1
TANF beneficiary in quarter	0.064	0.100	0.064
UI beneficiary in quarter	0.083	0.083	0.065
Felony conviction in quarter	0.012	0.012	0.014
Employed in quarter	0.92	0.93	0.53
Quarterly earnings (mean), cond. ($)	2533.9	2363.4	4030.9
Quarterly earnings (median), cond. ($)	1847.7	1756.5	3563.1
Newly hired in quarter	0.86	0.87	0.15
New job stints			
Total earnings at job (mean)	13563.9	12530.4	14470.1
Total earnings at job (median)	2771.0	2578.9	2476.6
Num. quarters with firm	3.50	3.42	3.54
Industry: retail	0.31	0.31	0.14
Industry: admin/support	0.33	0.33	0.27
Industry: food/accom	0.17	0.17	0.18
Applied for WOTC	1	1	0.21
Certified for WOTC	0.55	1	0.17
Pre-certified	0.012	0.015	
Starting wage (WOTC)	9.15	9.02	
Total hours at job (WOTC)	1483.4	1312.3	
Total hours: 0-120 (WOTC)	0.63	0.34	
Total hours: 120-400 (WOTC)	0.13	0.23	
Total hours: 400+	0.24	0.42	
WOTC subsidy ($)	608.5	1097.9	
Unique individual-quarters	640818	354829	10985951
Unique individuals	415828	234557	867916

Note: Table reports summary statistics for individuals in quarters during which: a firm applied for WOTC on their behalf (column 1), a firm received a WOTC certification on their behalf (column 2), or they had received SNAP in three of the past five months and were between the ages of 18-39. Data from 2005-2019.

Table 2.2: Effect of WOTC on individual outcomes, DiD analysis from SNAP expansion

Panel A: Stacked sample						
	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Treat x Post	−0.0016	0.0064	17.5775	0.0110***	8.7523	0.0001
	(0.0021)	(0.0042)	(24.5426)	(0.0006)	(8.0238)	(0.0007)
Dep var mean	0.1546	0.5206	1,934.3659	0.0180	877.7501	0.0154
Observations	594,524	594,524	594,524	594,524	594,524	594,524
Panel B: Age 25 cutoff sample						
	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Treat x post	−0.0001	0.0064	54.6391*	0.0132***	3.4922	−0.0001
	(0.0029)	(0.0053)	(30.6431)	(0.0010)	(10.0911)	(0.0009)
Dep var mean	0.1883	0.6046	2,147.7202	0.0282	949.8797	0.0163
Observations	365,651	365,651	365,651	365,651	365,651	365,651
Panel C: Age 40 cutoff sample						
	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Treat x post	−0.0040	0.0065	−41.5991	0.0075***	17.1513	0.0004
	(0.0030)	(0.0069)	(40.8417)	(0.0006)	(13.2134)	(0.0011)
Dep var mean	0.1012	0.3871	1,595.2563	0.0019	763.1060	0.0141
Observations	228,873	228,873	228,873	228,873	228,873	228,873

Note: Table reports difference-in-differences estimates of the effect of WOTC from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.2) on new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Controls include monthly age fixed effects and quarter-year by age sample fixed effects in Panel A. Controls include monthly age fixed effects and quarter-year fixed effects in Panels B and C. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.3: Effect of WOTC on individual outcomes, RD analysis

	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Panel A: Full sample						
RD_Estimate	−0.0007	−0.0002	−3.1037	−0.0097***	0.8321	0.0000
	(0.0010)	(0.0020)	(12.0372)	(0.0004)	(4.2601)	(0.0003)
Dep var mean	0.124	0.469	2,032.539	0.003	1,025.865	0.011
Bandwidth	2.858	3.082	4.272	1.585	3.539	4.627
Effective Obs	2,782,835	2,947,210	4,174,431	1,255,960	3,438,705	3,999,889
Panel B: Unemployed sample						
	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
RD_Estimate	0.0004	0.0008	4.4435	−0.0108***	−1.4682	−0.0001
	(0.0013)	(0.0015)	(4.9398)	(0.0005)	(5.8701)	(0.0004)
Dep var mean	0.126	0.151	333.487	0.003	1,135.288	0.014
Bandwidth	2.672	2.592	2.444	2.147	3.510	4.471
Effective Obs	1,602,980	1,552,567	1,450,598	1,022,585	2,109,348	2,362,483

Note: Table reports regression discontinuity estimates for employment, earnings, program utilization, and criminal activity. The running variable is distance between the individual's age and the age cutoff for WOTC eligibility (age 25 before 2007 and age 40 in 2007 and afterwards). Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.4: Effect of eWOTC on individual outcomes

	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
Treat x post	0.0032	−0.0009	−61.0683*	0.0038***	−3.2468	0.0002
	(0.0023)	(0.0052)	(34.6924)	(0.0007)	(8.0994)	(0.0006)
Dep var mean	0.1045	0.4312	2,073.8991	0.0039	764.3424	0.0092
Observations	458,986	458,986	458,986	458,986	458,986	458,986

Note: Table reports difference-in-difference estimates from the 2017 shift to electronic filing of WOTC applications among our SNAP sample for new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Controls include monthly age fixed effects and quarter-year fixed effects. Standard errors clustered at the individual level.

Table 2.5: Effect of WOTC on firm-level outcomes

	Panel A (40-41 vs 38-39)		Panel B (40-44 vs 35-39)	
	Diff. in hires	Diff. in WOTC certs.	Diff. in hires	Diff. in WOTC certs.
ATT	0.024	0.100***	0.026	0.310***
	(0.020)	(0.008)	(0.041)	(0.023)
Observations	51,884	51,884	51,884	51,884

Note: Table reports difference-in-difference estimates of the staggered firm-level adoption of WOTC on hiring and utilization of WOTC, following Callaway and Sant'Anna (2021). Panel A defines the outcome as the difference in the number of hires (WOTC certifications) who are SNAP beneficiaries aged 38 or 39 (i.e., eligible for WOTC) and SNAP beneficiaries aged 40 or 41 (i.e., just eligible for WOTC). Panel B defines the outcomes using SNAP recipients aged 35 to 39 as the eligible group and those aged 40 to 44 as the ineligible group. All regressions include firm and year fixed effects. Standard errors clustered at the firm level.

Table 2.6: Initial Job Applications for a Sample of WOTC-using Firms

Panel A: All Applications (N=69)	
Characteristic	(% of applications)
WOTC Screening Questions	13.04
EEOC Questions	42.03
SSN Required	18.84
Consent to Background/Drug Test	23.19
Submit Resume	44.93
Education History	53.62
Employment History	66.67
References	8.70
Referral	33.33
Personality Test	2.90
Panel B: Applications with WOTC Screening Questions (N=9)	
Characteristic	(% of applications)
WOTC questions direct to third-party site	44.44
WOTC questions stated as voluntary	88.89
WOTC questions ask about individual target group eligibility	66.67

Note: Table reports means of job application characteristics collected from a representative sample of online job applications for WOTC-using firms.

Chapter 3

Does Affirmative Action Impact Inter-Generational Mobility? Evidence from India

3.1 Abstract

I estimate the causal impact of an affirmative action policy, the implementation of job quotas for a socio-economically disadvantaged group in India, the Other Backward Castes (OBCs), in 1993, on their inter-generational education and occupation mobility. I use data drawn from the Indian Human Development Survey (2011-12), and employ a difference-in-differences strategy motivated by the fact that only OBCs of a school or college-going age in 1993 could potentially be impacted. Older OBCs would have already made human capital investment decisions that would be impossible or too costly to change in response to the policy. I find that the job quotas resulted in an increase in the absolute upward education mobility of OBC sons, likely reflecting heightened educational aspirations, as indicated by (1) an increased probability that their education is greater than their father's and (2) an increase in the years of education of sons born to illiterate fathers. On the other hand, I find that the policy resulted in an absolute downward occupation mobility of OBCs, as indicated by a decrease in the probability that a son born to a father employed in an unskilled or farming occupation is employed in a professional or skilled occupation, even though public-sector employment among OBC youth expanded.

3.2 Introduction

Do affirmative action policies help in improving Inter-Generational Mobility (IGM) of the communities targeted by such policies? I address this question in the context of the implementation of the Mandal Commission reforms, which were a series of reservation policies that introduced a 27% reservation of seats in central government jobs and colleges for the Other Backward Castes (OBCs), a historically disadvantaged group in India, in 1993 and 2006, respectively. In this paper, I focus only on the impact of the job quotas.

Inter-generational mobility, or the degree to which the socio-economic status of a generation persists, is often used as an indicator of the equality of economic opportunities in a society (e.g., Piketty (2000); Restuccia and Urrutia (2004)). Understanding and measuring this degree of persistence, as well as unravelling the respective contributions of the transmission of innate abilities, family background, and economic policy in generating this persistence, has been one of the most controversial issues, not only in actual political conflicts, but also in academic writings by social scientists (e.g. Lee and Seshadri (2019); Piketty (2000)). The ultimate aim of affirmative action policies is to ensure equal access of opportunities by all members of the society; hence, it is a natural question to study their effect on inter-generational mobility.

There are five motivations for studying the impact of affirmative action laws on inter-generational mobility in the Indian context. First, affirmative action policies in India are larger, more aggressive, and more salient than in most other countries (Khanna, 2020). Second, over time, reservation policies have increased in their scope, rather than decreased. Third, these policies are widely discussed in households and often dominate the public discussion regarding education policies in India. Public colleges and jobs, in which these affirmative action policies are implemented, are still considered very attractive and prestigious. Thus, I expect households to be well aware of these highly-publicized policies and take them into consideration when making human capital formation decisions for their children. Fourth, studies on IGM that look at causal impacts of policies based in India are rare. Fifth, studying inter-generational mobility against the background of the social architecture shaped by the caste system is even more interesting as the caste system served as a natural barrier to

IGM by restricting individuals to their traditional caste-based occupations and preventing them from marrying outside their caste community.[1] Each caste was traditionally linked to (usually) a single occupation (Cassan et al., 2021), and individuals were fated to follow that occupation by virtue of their birth.

Even today, caste as an institution plays a vital role in all stages of an individual's education and labor market outcomes in India, by determining the quality of school and college that they have access to, how they are treated within that school or college by their teachers and peers, whether they can benefit from affirmative action in access to colleges and government jobs, how they are perceived by potential employers and clients, and the informal caste-based networks they can rely on for job referrals and credit (Munshi, 2019). Figure 1 shows how the education of general category individuals is consistently higher than the education of OBC individuals, especially for those born in earlier birth cohorts. Similarly, Figure 2 portrays that general category individuals are more likely to be employed in professional or skilled occupations, which on average, offer a higher pay than other occupations. Caste also plays an important role in the political sphere, and indeed any discussion about caste, especially caste-based affirmative action policies, is a politically sensitive topic. Changes in caste-based affirmative action policies are often met by large protests and political parties often indulge in vote-bank politics, by promising to target public resources to certain castes, even at the local level. Thus, the caste system makes India an important setting to understand the inter-generational persistence in socio-economic outcomes and also to isolate the impact of group-based affirmative action policies on inter-generational mobility.

Theoretically, the job quotas could influence the education and occupation choice of the OBCs through the following two mechanisms. First, the job quotas could have a direct positive impact on the occupation choices of the OBCs as a result of expanded opportunities in the public sector. However, these opportunities can be availed only by acquiring at least the minimum level of education required for these jobs. At the same time, private sector employers might start perceiving OBCs who fail to secure a government job as inferior candidates, and thereby reduce their probability of employment (Coate and Loury, 1993).

[1] This is the view taken by most historians, but some historians like Chandavarkar (2003) and Rudner (1996) theorize that mobility is possible in a caste-based society, but this mobility means the mobility of the entire group, rather than of an individual.

Negative stereotypes about the ability of minority candidates would be exacerbated and could lead to more discrimination by employers. Second, the job quotas could alter the human-capital-formation decisions of OBCs due to a change in the perceived returns to education. A higher investment in education could, in turn, lead to more opportunities in the private sector. However, such policies can also have a negative impact by reducing the competition for public sector jobs, and thus, reducing the incentive to invest in human capital formation.

There are two strands of literature that are related to my work: (i) papers that look at the impact of affirmative action, for example, on the education and labour market outcomes (the probability of being employed, the probability of being salaried, etc.) of the targeted minorities and (ii) papers that focus on empirically estimating measures of IGM. To the best of my knowledge, my paper is the first one to attempt to connect these two strands of literature by trying to estimate the causal impact of an affirmative action policy on intergenerational mobility.

Extensive literature exists on affirmative action and its impact on education and labour market related outcomes. Khanna (2020) examines the effect of the same policy that I study, i.e., the implementation of the job quotas for the OBCs, and finds that OBC students are incentivized to study in school for another 0.8 years on average. The younger the OBC student is, the higher is the expected increase in their educational attainment as they would have more time to adjust their human capital formation decisions in response to changes in future prospects. Prakash (2020) uses exogenous variation in the the updating of state-level job quotas for the Scheduled Castes (SCs) and the Scheduled Tribes (STs), two other disadvantaged sections in India, and concludes that a one percentage point increase in the job quotas for the SCs raises their probability of being salaried by 0.6%. However, no such effect is observed for the STs. I study the same policy and use the same methodology as in Khanna (2020), but my main objective is to look at measures of IGM. One of the most common criticisms leveled against affirmative action policies in India is "elite capturing", and thus, it is important to examine if the benefits of the job quotas are concentrated among members who are born to relatively better-off parents.

The empirical literature on IGM aims to measure the degree to which a parent's socioeconomic status determines their child's socioeconomic status (Chetty et al., 2014). Because opportunities and socioeconomic status are hard to measure, most of the literature on IGM is set in developed countries and considers income to be the most informative measure of socioeconomic status (e.g. Chetty et al. (2014); Solon (1999)). However, IGM studies set in developing countries (or even historical periods set in developed countries[2] tend to avoid using income and instead use education (Alesina et al., 2021; Asher et al., 2018; Azam and Bhatt, 2015; Hnatkovska et al., 2013) or occupational choice (Azam, 2013; Hnatkovska et al., 2013; Iversen et al., 2017) as the relevant proxy of socioeconomic status for the following reasons.

First, income in developing countries is subject to significant measurement errors. Availability of long panel data is rare, and studies typically rely on cross-sectional data that record an individual's income in the year in which the survey was conducted. This measure of income tends to be a poor estimate of permanent lifetime income as it is affected by transitory income shocks, which are especially prevalent in economies predominated by the agrarian and informal sectors. Analysis of income mobility requires that we compare the outcomes of children to their parents when they are at similar ages in their life-cycle because measuring children's income at early ages can positively bias inter-generational mobility in lifetime income as children with high lifetime incomes have steeper earnings profiles when they are young (Grawe, 2006; Zimmerman, 1992). In contrast, education and occupation choices rarely change in adulthood and thus, do not suffer from this life-cycle bias.[3] Education and occupation also tend to be precisely reported whereas income is often under-reported.

Second, many households in developing countries engage in joint production activities, like cultivation and self-employment in a family-owned business, which sometimes makes it impossible to divide the household income among the family members.[4] Education and occupational choice, in contrast, can directly be attributed to an individual. Third, infor-

[2]Examples of studies looking at educational inter-generational mobility in historical times in developed countries include Derenoncourt (2022) and Card et al. (2022).

[3]I still try to control for the life-cycle bias by adding age and the square of age of the son and the father as covariates.

[4]One possible way of allocating household income among household members could be to use adult equivalent scales. The survey data set that I use, the IHDS-II, reports the number of hours worked by each household member in joint production activities, and hence, another method could be to divide the household income in accordance with their time contribution. However, the other two issues in using income as a measure of socio-economic status still remain.

mation on education and occupation of parents might be available even if they are dead. However, education and occupation are not perfect measures of socio-economic status. The emphasis on studying educational mobility in developing countries is driven by the belief that education opens up better opportunities in the future, but the returns to education vary for different social groups. The empirical literature at best compares the trends of the estimates of IGM or explains how the estimates could be different for different social groups, without examining potential reasons that could explain such differences. My contribution to the literature is to investigate the role of one possible reason (i.e. affirmative action) in explaining the differences in IGM among different social groups.

To study the effect of affirmative action implemented in the form of job quotas on educational IGM, I look at the following indicators of absolute upward education mobility: (1) probability that a child acquires higher years of education as compared to their parent, (2) years of education of children born to illiterate parents, and (3) probability that a child born to an illiterate parent is literate, Similarly, to look at absolute downward education mobility, I use the following measures: (1) years of education of children born to at-least-high-school-graduate parents and (2) probability that a child born to an at-least-high-school-graduate parent does not complete high school. I restrict my analysis to only son-father pairs because tracking parental information for married women is challenging in any Indian household survey, and the labor force participation rate for women is low.

Occupations, unlike education, can not be easily ranked as there is some subjectivity in what a "better" occupation is.[5] Thus, I classify the occupations of individuals as (1) professional, (2) skilled, (3) unskilled, and (4) farmers. To study absolute upward occupation mobility, I use the probability that a son who is born to a father employed in an unskilled or farming occupation is employed in a professional or skilled occupation. Similarly, to study absolute downward occupation mobility, I use the probability that a son who is born to a father employed in a professional or skilled occupation is employed as a farmer or in an unskilled job.

I use data drawn from the second round of the Indian Human Development Survey

[5] One way to rank occupations could be to use the average income from those occupations. Another way, which was popularly used by American sociologists in the 1960s, is to rank occupations based on occupation prestige. See for example, Nakao and Treas (1994).

(IHDS), which was conducted in 2011-12, and construct a data set of matched son-father pairs. I then restrict my sample to OBC and general-category individuals who have completed their education at the time of the survey. Next, I exploit the plausibly exogeneous nature of the implementation of the job quotas and employ a difference-in-differences strategy to estimate the causal impact of the job quotas for OBCs on inter-generational education and occupation mobility. I do so by comparing the outcomes of OBC sons belonging to a particular subsample based on father's education or occupation, to their counterparts in the general category, for individuals who are "old" at the time of the policy enactment and those who are "young". Only OBCs of a school or college-going age in 1993 could potentially be impacted by the job quotas as older OBCs would have already made human capital investment decisions that would be impossible or too costly to change in response to the policy.

I find that the job quotas resulted in an increase in the absolute upward mobility of OBC sons, as indicated by (1) an increased probability that their education is greater than their father's, and (2) an increase in the years of education of sons born to illiterate fathers. On the other hand, I find that while the policy resulted in expanded public-sector employment for OBCs, there is a slight absolute downward occupation mobility of OBCs, as indicated by a decrease in the probability that a son born to a father employed in an unskilled or farming occupation is employed in a professional or skilled occupation. Moreover, for almost all of these results, I find that the youngest cohorts (ages 8–12 in 1993) respond most strongly, followed by those aged 13–17, with the effects weakest among those aged 18–22. Overall, the evidence suggests that while OBC sons acquired more education in response to perceived future opportunities, but the downward occupation mobility could perhaps be understood by the limited opportunities or exacerbation of negative attitudes of private employers regarding OBCs in response to the quotas.

The rest of the paper proceeds as follows. Section 3 discusses the institutional background of reservation policies in India, and in particular, the implementation of the Mandal Commission reforms. Section 4 provides a description of the data source, the construction of the data set, and the variables used in the econometric analysis. Section 5 details the empirical strategy employed to study the impact of the job quotas on inter-generational mobility. Section 6 reports the estimation results and interprets the main findings. Section 7 concludes and discusses some ways in which the analysis could be extended in the future.

3.3 Institutional Background

I divide this section into four subsections. In the first subsection, I explain what the caste system is. The second subsection details the history of reservation policies in India and the third subsection discusses the particular reservation policy that I am interested in, i.e. the implementation of the Mandal Commission reforms. The fourth subsection presents some facts on the public sector in India.

3.3.1 Caste System

The majority of reservation policies in India are based on caste and take the form of quotas in educational institutions, jobs, and legislatures, which means that only the eligible groups for the reservation policies can apply for a specified percentage of seats. Thus, before proceeding to discuss the reservation policies in India, it is crucial to provide a little background on what the caste system is.

The caste system was a hereditary and hierarchical way of organization of society and division of labour in ancient India that was developed between 1500 to 500 BCE (Munshi, 2019). The caste system is characterized by the following main features. Every Hindu individual was (and for the most part still is) born into a *jati*, or caste, which in turn belongs to a *varna*, though there is sometimes some ambiguity about which *jati* belonged to which *varna*. The *varna* system stratified the Hindu society into four hierarchical classes (called *varnas*), with the *Brahmins*, who were priests and scholars, at the top of the social hierarchy. Next were the *Kshatriyas*, or the warriors and rulers. They were followed by the *Vaishyas*, or merchants, and at the bottom of these four caste categories were the *Shudras*, who were usually laborers, peasants, artisans, and servants. Outside of the *varna* system were a large sub-population of *Dalits* or untouchables who were excluded from much of the Indian village life. The other sub-population that was not a part of the *varna* system was the tribals. Thus, the *varna* system defined a broad social structure within which the thousands of castes and subcastes were placed. There are about 2,000-4,000 castes in India, and many more subcastes. It is important to realize that it is the *jatis* that define and form the nuts and bolts of the caste system, rather than the better-known *varna* system (Osborne, 2001). Each *jati* was linked to a particular occupation and tended to be confined to particular regions. Each *jati* was obligated to perform certain services tied to their hereditary occupations and could draw upon the services that other *jatis* were ordained to follow. The resulting *jajmani* system, which was the complex system of obligations among *jatis*, mimicked exchange in a conventional monetary economy (Osborne, 2001).

The caste system was very rigid, with strict restrictions on inter-caste dining and especially inter-caste marriage, to prevent mixing of castes. A notion of purity and impurity emerged, with the *Brahmins* considered the most "pure", and purity declining successively with *Kshatriyas*, *Vaishyas*, *Shudhras*, and the untouchables. A touch from a lower caste was considered to"pollute" the upper caste. The caste system, with all of its taboos and restrictions, was authenticated by religion, and the justification behind this hierarchy was tied to the Hindu belief in rebirth, wherein a person who was born as a *Brahmin* was considered to be able to do so by virtue of doing good deeds in their past life.

Of course, the system was unfair (and inefficient too) as one's occupation, and thus, fate was determined at birth, permitting no occupational or socio-economic mobility and resulting in persistent inequality across generations. This is the main feature that distinguishes caste from class, as the latter is not hereditary, at least in principle. Over time, the caste system lost its influence due to the influx of other cultures and religions, and it was believed to be fairly flexible by the 18^{th} century. Harder boundaries were again set by the British, who wanted a way of classifying a large and diverse Indian population, and they made caste India's defining social feature. They also practised discriminatory policies, such as allowing only Christians and upper-caste Hindus to be eligible for administrative positions in the British Raj.

Today, caste has became linked to religions apart from Hinduism as well. Many communities that converted from Hinduism to Islam, Christianity, or other religions, still maintained their caste and subcaste affiliations.[6] Caste identities remain strong, and untouchability

[6]Thus, non-Hindus could also be eligible for caste-based reservation policies.

and prohibition on inter-caste dining is still practised in many remote areas. Field-based ethnographic studies in villages present mixed evidence regarding the hereditary relationship between occupations and caste (Jodhka, 2004; Mayer, 1996; Mendelsohn, 1993), thus, highlighting the need to use large scale survey data to explore if the relationship between caste and occupational segregation has weakened, and if so, how much can be attributed to affirmative action policies. Inter-caste marriages remain low, at around 5.8%[7] according to the 2011 Census and this rate has remained virtually unchanged over the past four decades.

3.3.2 History of Reservation Policies in India

Caste-based discrimination, especially the practise of untouchability, was banned after independence in 1947 in hope that the caste system would cease to be a prominent part of Indian society. However, the post-independence government decided to continue the British policy of reserving seats, a policy, ironically, that had been opposed pre-independence (Osborne, 2001). The difference, though, was that now, affirmative action policies were introduced to correct inequalities arising from the caste system and explicitly target the historically disadvantaged castes, whereas earlier, the policies favoured upper-caste representation.

Thus, for the purpose of caste-based reservation laws, there are 4 distinct caste categories in society today: (i) Forward Castes or General Category, (ii) Scheduled Castes (SCs), (iii) Scheduled Tribes (STs), and (iv) Other Backward Castes (OBCs). This categorization is different from the original *varna* system of organizing the castes. However, there is some general overlap, for example, the castes comprising the SCs are mostly the castes classified as untouchables in the original *varna* system categorization. Similarly, the *jatis* comprising the STs are mostly the aboriginal tribal communities that existed outside the *varna* system. *Brahmins* are generally a part of the forward castes, while the *jatis* that belongs to the *Kshatriyas*, *Vaishyas*, or *Shudras* are generally a part of either the forward castes, or the OBCs, depending on the region.

The forward castes are considered the most advantaged sections of the society, and the SCs and the STs are considered the most disadvantaged. From a policymaker's perspective,

[7]Ray et al. (2020) study the patterns of inter-caste marriage and find that the rate is not statistically different for rural and urban areas, and education of the spouses themselves plays no role, but education of the mother of the groom does.

there are at least three factors to consider while forming reservation policies: (i) the category of individuals to target i.e. which castes belong to these four caste categories, (ii) the institutions in which to introduce reservations (schools/colleges/jobs/legislatures and also the the jurisdiction level– central/state/local/private), and (iii) the percentage of seats to reserve. The government comes up with a list of castes and sub-castes to classify as SCs, STs, and OBCs, and by default, all other castes and sub-castes that do not belong to either of these three lists are a part of the Forward Castes.[8]

The first round of reservation policies were implemented in 1950 and were targeted at the most disadvantaged social groups, the SCs, and the STs. The SCs and the STs occupy a 15% and 7.5% reservation of seats, respectively, in central government colleges and jobs. In 1993 and 2006 respectively, the OBCs also started occupying a 27% reservation in government jobs and colleges, which I will discuss in detail in the next subsection. The OBCs, however, do not enjoy reservations in central, state, and local legislatures, as the SCs and the STs do.

To put these quotas into perspective, it would be helpful to know the share of population of each of these four caste categories. The counting of the population of different castes in India was started by the British in 1872,[9] but this practise was discontinued after the Census of 1931. Caste data in the 1941 Census was collected, but not published. The Post-independence decennial Censuses of India, first conducted in 1951, record if an individual belongs to the SC or the ST caste group, but do not record if an individual belongs to the OBC group. According to the 2011 Census, the share of SCs and STs in India's population in 2011 was 16.6% and 8.6%, respectively. This means that the General and the OBC caste group together constituted 74.8% of India's population in 2011. The Mandal Commission that was established to identify OBC castes estimated their popluation share to be 52% in the early 1980s, but there aren't any official figures available since then. While the quotas for the SCs and the STs are somewhat in tune with their population shares, 27% was recommended

[8]These lists are different for different states, thus it is possible, for example, that the same caste could be considered OBC in one state and SC in another state.

[9]In 1872, 3,208 castes were identified. The castes in colonial India weren't divided into the four caste groups (SC/ST/OBC/General) that they are divided into now, hence there isn't any information available in the colonial censuses on the share of population belonging to these different caste groups. It might be possible to aggregate the 1872 population of all the castes that belong to a particular caste category in 2021 (the list of OBC/SC/ST castes is publicly available) to arrive at these shares. However, I have not carried out this exercise and I couldn't find any other paper that has.

as the quota for the OBCs by the Mandal Commission, despite their population share being estimated by them to be 52%, owing to a Supreme-court mandated constraint that the total quantum of reservation could not exceed 50%. Nationally representative surveys, like the National Sample Survey (NSS), estimate their number to be around 40%, but the reluctance to officially count the OBCs could stem from the government's fear that the official count could be above 52%, thereby spurring demands for more quotas. The first ever Socio-Economic Caste Census (SECC) was conducted in 2011, which reported the caste of every respondent, but the results of the caste count have not yet been released.[10]

The preceding discussion on the quota shares for different caste groups is at the central level. Besides central government jobs and colleges, one could also apply to colleges owned by state-level governments or get a job in state-level institutions. Different state governments have their own policies regarding reservation in state government jobs and colleges. Each state chooses a different reservation quotas for the SCs, STs, and the OBCs, and these quotas could be different from the central quotas of 17%, 7.5%, and 27%, respectively. [11] Khanna (2020) provides information on the state-level variation in the quotas applicable to different caste categories in the year 1995, though there is a slight possibility that these numbers may have changed over time.

One concern could be that the list of SCs, STs, and OBCs is in itself determined by factors like political lobbying wherein castes with higher resources and political clout can pressure the government to make them eligible for quotas. Recent anecdotal evidence suggests that it may not be that likely. One caste community, the Jats, led a series of violent protests in February 2016 to seek inclusion in the OBC list of the central government and the state governments of Punjab and Haryana. The state governments ceded and categorized them as OBCs, but the decision was overruled by the court within a couple of days.

Besides caste-based quotas, affirmative action policies are also targeted towards women and low-income general category caste members. In 1992, a constitutional amendment man-

[10]This is because of the following two reasons. First, the government has faced logistical difficulties in trying to collate caste information. Around 330 million households were surveyed, and 4.3 million entries were recorded, which clearly exceeds the actual number of castes in India. There are differences in the interpretation of the question "What is your caste?" by different individuals as caste is a complex social structure. Second, the ruling party could be deliberately delaying releasing the results so that they don't have to deal with the resulting political fallout.

[11]Thus, in my analysis in Section 3.5.1, I control for State-Caste fixed effects.

dated one third of local government head positions to be randomly reserved for women. In 2019, 10% of seats in government colleges and jobs were reserved for low-income general category caste members. There have also been demands to have quotas in place separately for Muslims. The issue of sub-categorization of OBC quotas has also gained policy attraction, with more than 13 Indian states instituting a quota-within-quota system.

3.3.3 Mandal Commission Reforms

The identification and implementation of reservation quotas for the OBCs has been a long drawn-out process. Adhering to Article 340 in the Constitution, which obligates the government to promote the interests and welfare of "backward classes", the First Backward Classes Commission was set up on January 29, 1953 under the chairmanship of Kaka Kalelkar. Its objective was to identify sections of society that were socially and educationally backward, in addition to the SCs and the STs, and make recommendations to improve their condition. The commission submitted their report in 1955 and identified 2,399 backward castes, out of which 837 were deemed to be the "most backward". However, the report was not accepted by the government as caste was used as the sole criteria to identify social and educational backwardness, and the government feared severe backlash. The second Backward Classes Commission was established under the leadership of B.P. Mandal (and hence it is popularly called the Mandal Commission) on January 1, 1979, with the same objective. The Mandal commission came out with their report in 1980, wherein they created a list of castes that would be OBCs based on an index of backwardness that captured their socio-economic status at that time and also recommended the percentage of seats to be reserved for them in educational institutions and jobs. The index was based on four social indicators, each carrying a score of three points; three educational indicators, each carrying a score of two points; and four economic indicators, each carrying a score of one point. Thus, the total score could be at maximum 22. These 11 indicators were applied to each caste surveyed by the commission, and castes that scored less than 11 points were deemed backward. 3,473 backward caste (or subcaste) groups were identified.

The contents of the Mandal Commission report were of a politically contentious nature, and the central government did not act upon it for nearly 10 years. They first announced their intent to implement it in 1989, but were immediately faced by (often violent) protests by upper-caste students. In 1991, they again announced that the recommendations will be implemented, but the constitutional validity of the new law was challenged in court. On 16 November 1992, the Supreme Court ruled in favour of the central government and the first stage of the reforms was implemented in 1993. Thus, even though the recommendations were around for a long time, it can be argued that the implementation of these reservation policies was exogeneous from the perspective of the households, though there is some degree of endogeniety in the sense that it was implemented only after a change in the political party in power. Reservation policies take an important position in electoral promises; so, it would be the same households who vote for the parties, keeping in mind the expected policy changes in the future.

3.3.4 Public Sector in India and Implementation of Job Quotas

The public sector in India still remains a major employer, though the size of the public sector is shrinking over time. The formal/organized sector in India employed around 20% of the workforce in 2018, out of which around 58% comes from the public sector. However, the central government constitutes only 14% of the public sector employment, and the rest comes from state governments. Employment in the public sector is still considered very attractive, as it is seen as a symbol of status and power, provides a competitive salary and perks, and offers job security. This is true especially at the lower end of the skill distribution. For example, the unconditional annual wage premium for government jobs relative to similar formal sector jobs was $465 ($485 for OBCs) in 2000 (Prakash, 2020). Of course, the number of job openings as a percentage of the labour force is low. Competition for these jobs is tough, with one instance in 2018 where 23 million people applied for about 100,000 low-skilled posts advertised by the Railway Recruitment Board, such as porters.

Jobs in the central government of India are classified in four categories (group A, B, C, D)[12] corresponding to the salary, status, qualifications, and the nature of responsibilities attached to the jobs. Group A posts are the most well-paid, carry higher administrative

[12] The quotas are applied to each of the four categories separately.

and executive responsibilities and include senior management positions in the ministries/departments and field organisations. Group B jobs constitute middle management. Group A and B posts together made up 11.5% of the total central government jobs and require a higher secondary or a college degree. Group C employees perform some supervisory and operative tasks. These jobs comprise 58% of the jobs and require completing secondary or middle school. Group D employees receive the lowest salary and carry out low-skilled routine tasks such as clerical work. These jobs encompass 30.5% of the jobs and require completing primary school or being literate. Thus, the incentive effect of the job quotas could matter for acquiring each of these minimum education qualifications.

As mentioned before, hiring in the public sector as well as admissions to government universities are subject to caste-based quotas. In practise, the hiring takes place as follows. First, the government announces vacancies for a particular post, the minimum educational qualification and the maximum applicant age required for that post. An applicant can indicate that they intend to apply for a reserved seat, by showing proof of their SC/ST/OBC status.[13] These caste certificates are issued by the sub-district administration. Recruitment is determined based on written competitive exams[14] and the government comes out with a score cutoff for each caste separately. The score cutoffs are usually such that the top 50% of the highest scoring candidates in the general category can be hired, the top 27% OBC candidates can be hired, and so on, or until there is no available candidate that meets a "minimum quality" threshold. The cutoff for the general category is mostly higher than the cutoffs for other caste categories. "Minimum quality" is determined according to the nature of the vacancy, but is ambiguous. Reserved seats that are unfilled could be filled by lowering the cutoff for general category individuals. Thus, in practise, it is possible that more than 50% of the seats go to the general category, or some seats are unfilled even after determining multiple cut-offs. Quotas are also implemented in promotions.

Data on actual shares of different caste groups in government jobs are scant, but some

[13]SC/ST/OBC individuals could also apply as a general category individual by choosing not to show their caste certificate. This could be due to the fear that their peers or superiors in their job could discriminate against them if they reveal they are from a lower caste.

[14]For some positions, there could be more than one round of written exams. There could also be an interview as the final stage of the hiring process. In such cases, the final score of a candidate is determined as a weighted average of their scores in all of the hiring stages. The weights are generally such that more importance is given to written exams.

information is available through the replies of government officials to such concerns raised in the parliament or in their response to Right to Information (RTI) requests made by the public. The representation of SCs, STs, and OBCs in central government jobs stood at 17.49%, 8.47% and 21.57%, respectively, as on January 1, 2016. The OBC representation has increased over time; on Jan 1, 2012, the share of OBC employees in central government jobs was 16.55%.[15] Also, to give an idea of the OBC representation before the quota implementation, the Mandal Commission report estimated that their share was 12.55% in 1980 (4.69% in Class I jobs, 10.63% in Class II jobs, and 24.40% in Class III jobs).[16] The actual representation of the OBCs in government jobs falls short of the quota available to them, but this shouldn't be suggestive of the policy being ineffective. The effect of the intention of the government to fill seats such that the OBC share is 27% should be seen more from an Intent-to-Treat framework. The very availability of the quotas could change how much education young children acquire in response to perceived future return, in an aspirational sense.

3.4 Data

I divide this section into six subsections. The first subsection discusses the source of the underlying data and the advantages and disadvantages of using it. The second subsection describes the cprocess of linking sons to fathers. he choice of the control caste groups and the treatment cohorts is explained in the third subsection. The fourth subsection details the sample restriction criteria imposed. The fifth subsection describes the variables used in the econometric analysis. Finally, the sixth subsection presents summary statistics.

3.4.1 Data Source

To study the impact of the Mandal Commission reforms on inter-generational education and occupation mobility, I must have, at the very least, data on caste group membership,

[15]These aggregate numbers mask heterogeneity in the actual type of jobs secured by different caste groups. For example, lower caste groups are typically better represented in low-skilled jobs than in high-skilled jobs.

[16]The central government earlier used Class I/II/III to classify jobs. Class I jobs were the most sought-after and prestigious jobs, like civil servants, professors, etc. while Class III jobs were unkilled jobs like clerks.

age, years of education, and occupation for children as well as their parents. The data used in this study come from the second round of the Indian Human Development Survey (IHDS-II), which was conducted by the University of Maryland and the National Council of Applied Economic Research in 2011-12. IHDS is a multi-topic, nationally representative survey[17] and encompasses questions on a wide range of socioeconomic and demographic characteristics such as health, education, employment, economic status, marriage, fertility, gender relations, and social capital. Children aged 8-11 completed standard reading, writing and arithmetic tests, whose questions were the same for all the surveyed children. Additional village level data is also available on the status of schools, medical facilities, and other public infrastructure. IHDS-I surveyed 41,554 households in 1,503 villages and 971 urban neighborhoods across India whereas IHDS-II surveyed 42,152 households (in the same villages and urban neighborhoods), with 85 per cent of the households from IHDS-I being resurveyed.

There are six advantages of using data drawn from IHDS-II. First, the IHDS collects information on the education and occupation of the father of the majority of male respondents, even if those fathers have died or are not considered part of the same household.[18] Information on deceased fathers will help in understanding mobility of older cohorts. Second, it contains information on the education and occupation of non-resident children and non-resident parents. Most of the literature studying inter-generational mobility in India use data sets like the National Sample Survey (NSS) (for example, in Hnatkovska et al. (2013)) and National Family Health Survey (NFHS), which have information only on the residents of a household. Neglecting the non-residents might lead to negatively biased estimates, as individuals who experience upward mobility are also less likely to co-reside with their parents.[19] Third, the IHDS is a panel data set which allows me to recover the occupational

[17]IHDS covered all the states and union territories of India except Andaman and Nicobar and Lakshadweep islands. These two union territories form less than 0.05% of India's total population according to the 2011 Census of India.

[18]The household file in IHDS specifically asks the occupation and education of the father of a male household head. If the household head is female, they instead ask the occupation and education of her father-in-law.

[19]Asher et al. (2018) find that estimating IGM using only co-resident son-father pairs leads to a bias that increases in the age of the son. Azam and Bhatt (2015) demonstrate that the estimated regression coefficient in a regression of son's educational attainment on father's educational attainment restricting the data based on the co-residence criterion is 17% lower than the estimate based on the full sample. Other papers like Hnatkovska et al. (2013) also acknowledge that using only co-resident son-father pairs leads to a selection bias but do not show the possible extent of this bias.

status of some individuals who were retired in 2011-2012 (i.e. in IHDS-II, the data set that I use), but not in 2004-05 (IHDS-I). Fourth, the IHDS identifies the mother's and father's ID of an individual, if their parents are part of the same household. Data sets like NFHS and NSS do not provide an ID for the parents and children can only be indirectly linked to their parents using information from the relationship-to-the-household-head field. This, of course, is possible only in non-ambiguous cases, which might result in some information loss as many people in India still live in joint-families.[20] Fifth, it contains data on the actual number of years of education rather than levels of schooling, which helps me in avoiding bunching in schooling distribution as a result of the imputation of years of education from the level (Azam and Bhatt, 2015). Sixth, the IHDS provides very detailed information on a variety of household and village level characteristics, which allows for a very rich set of potential covariates.

Before proceeding, it is also important to acknowledge the limitations of IHDS-II for the purpose of this study. IHDS-II has a limited sample size as compared to data sets like the NSS and the NFHS, which are repeated cross-section surveys that have the advantages of a larger sample size and a higher number of rounds. The relatively small sample size of IHDS-II proves to be a limitation in some of my econometric analysis where in I subset the data by some variable indicating the father's education or occupation.

3.4.2 Linking Sons to Fathers

I limit my analysis to only son-father pairs due to the following reasons. First, the organization of society in India is patrilineal i.e. the bride moves into the groom's family post marriage. The parents of a woman who enters a household after marriage are, thus, not considered members of the household, and therefore it would be impossible to track the parental information for married women. Married daughters are also not considered a part of the household, and their information is not included in the non-resident file of IHDS. Second, the IHDS only records the education and occupational choice of fathers of male household heads, and not their mothers. Male household heads with a dead mother form a

[20]For example, Hnatkovska et al. (2013) are able to identify father's education for less than 15% of adult males surveyed in the NSS. In contrast, I am able to identify father's education for more than 95% of adult males in IHDS-II by using the father's ID and information on the education of the household head variable.

significant portion of my sample, and thus I can't ignore them while possibly trying to link sons with their mothers. Third, I restrict my sample to men because I want to study the impact of affirmative action on occupational mobility, but the female labor force participation in India still remains very low, at around 20% in 2019.

I construct matched son–father pairs using the household, individual (resident), and non-resident files of the IHDS. In the individual file, every resident is assigned an ID and asked their age, gender, occupation, earnings, relationship to the household head, and the IDs of their mother and father (if the parents are alive, regardless of whether they live in the household). The household file supplements this information by reporting the occupation and education of the father of a male household head, or of the husband's father if the head is female, even if that father is deceased or no longer part of the household. The non-resident file records the age, sex, occupation, and education of non-residents, the ID of the resident to whom they are related, and the type of relationship (e.g., son/daughter, father/mother, sibling). This structure allows me to link resident sons to resident or deceased fathers, as well as to non-resident fathers, and to link non-resident sons to resident fathers. In total, I identify 113,438 son–father pairs: 99,071 cases of resident sons linked to resident or deceased fathers, 2,991 cases of resident sons linked to non-resident fathers, and 11,376 cases of non-resident sons linked to resident fathers (Panel (A) of Table 3.1).

3.4.3 Control Caste Group and Treatment Cohorts

I use a difference-in-differences framework to estimate the causal impact of the implementation of the job quotas on the inter-generational education and occupation mobility. To identify the causal effects, I leverage variation in caste group affiliation and age. As discussed in Section 3.3.3, only the OBCs were eligible for the job quotas implemented in 1993. Thus, the OBCs are the obvious choice for the treatment group. This means that there are three potential candidates for the choice of the control group: (i) the general category/forward castes/upper-castes, (ii) the SCs, and (iii) the STs. Following the previous literature (Khanna, 2020), I keep only the general category as the control group. Khanna (2020) demonstrates that the impact of the job quotas for OBCs on their educational attainment was higher for each successive birth cohort affected by the policy change, and hence, it is not unreasonable to rule out dynamic treatment effects of the quotas for the SCs and the STs, especially since their quotas started in the 1950s and have covered more than one generation.

Another concern could be that caste-based reservation in legislatures could also have effects on policy outcomes through changes in the allocation of public goods. Pande (2003) finds that increasing caste-based minority representation in state governments increases transfers to minorities and results in an increased spending on welfare programs and share of seats reserved in state government jobs. Bardhan et al. (2010) find that having a SC/ST village head leads to redistribution of employment program benefits away from non-SC/ST landless households toward SC/ST households. In 1993, the year in which job quotas for OBCs were implemented, gender-based quotas were also introduced in local government bodies, mandating that one third of seats be reserved for women. This new policy, interacted with the pre-existing legislative quotas for the SCs and the STs, could thus change their socio-economic outcomes.

The source of the second difference in the difference-in-differences analysis comes from variation in age. The newly-implemented job quotas for OBCs could only be availed by individuals who were not too old at the time of the policy change. In Section 3.2, I explained that there are two possible mechanisms by which job quotas could affect IGM : (i) change in human-capital-formation decisions due to a change in the perceived returns to education and (ii) change in the probability of employment in certain occupations, even if there is no additional investment in education, as a direct result of expanded opportunities in the public sector and an indirect result from the change in how private employers perceive them. Government jobs have a maximum age cutoff and require some minimum educational qualifications (discussed in more detail in Section 3.3.4). It is also uncommon for older individuals to (i) acquire more education after they enter the labour market[21], even though educational institutions do not have a maximum age cutoff and (ii) drastically change their occupational choice, especially when occupation is considered in broader categories. Thus, if OBC individuals are young enough in the years in which the job quotas exist (i.e. of age 21 or

[21] E.g., according to data in in IHDS-II, in 2012, less than 0.5% of the individuals above the age of 25 were enrolled in any kind of educational institute.

younger–the age at which one typically graduates from college), they could have had time to change their education and occupational choice decisions in response to the implementation of the quotas. Individuals above the age of 21 will find it difficult to benefit from the job quotas as they would have already taken many human-capital-formation decisions that might now be too costly to alter. For the same reason, I drop individuals who were younger than 22 in 2006, as 2006 was the year in which quotas for the OBCs in central-government universities were introduced. Thus, I am able to capture the sole effect of the job quotas.

3.4.4 Sample Restrictions

From the constructed sample of son–father pairs, I impose several restrictions (Panel (B) of Table 3.1). First, I restrict the analysis to individuals belonging to either the General category or the OBC category, treating OBC members as the treatment group and General category members as the control group, following the precedent in the literature (Khanna, 2020). Second, I drop cases where the son's age (or birth year) is missing, since I need this information to classify individuals into "young" or "old" cohorts at the time of the 1993 introduction of OBC job quotas. I also drop cases where the recorded father's age is less than 10 years older than the son, as such observations are unlikely to be accurate. Third, I restrict the sample to birth cohorts 1951–1985. This ensures that all individuals were at least 21 years old in 2006, when OBC quotas in public educational institutions were introduced. Excluding younger cohorts avoids conflating the effects of educational quotas with those of job quotas, as individuals under 21 in 2006 could have adjusted their educational and occupational decisions in response to the education quotas. After applying these restrictions, the sample consists of 32,650 son–father pairs: 6,934 General (old), 9,086 OBC (old), 6,811 General (young), and 9,819 OBC (young).

I then impose my final set of restrictions (Panel (C) of Table 3.1). First, I exclude a few sons who were enrolled in any educational institution at the time the survey was conducted (i.e., in 2012). This is done to avoid the possible right-censoring of education data by the inclusion of sons who had not yet completed their education. Including such individuals could bias the estimates of intergenerational educational mobility downwards. Second, I sequentially drop cases with missing education—first for sons, then for fathers—so that only pairs with complete information on both are retained. The resulting sample consists of 26,467 son–father pairs for studying educational mobility: 5,919 General (old), 7,882 OBC (old), 5,137 General (young), and 7,529 OBC (young). I am able to preserve about 81% of the pairs I started out with (26,467/32,650). Father's education is recovered in 67% cases (17,889 of these 26,467 pairs) using information on the father of a male household head, even if the father was deceased.

Next, for studying occupational mobility, I take these 26,467 son–father pairs and further sequentially drop cases with missing occupation—first for sons, then for fathers. After imposing this additional restriction, I am left with 22,672 son–father pairs for studying occupational mobility: 5,191 General (old), 7,169 OBC (old), 4,046 General (young), and 6,356 OBC (young).[22]

3.4.5 Variables

Since I am interested in studying education and occupation IGM, I first describe the education and occupation related variables. I then proceed to describe other variables that serve as the main independent variables or as covariates.

3.4.5.1 Years of Education

IHDS-II reports the number of completed years of education of residents in the individual file, non-residents in the non-resident file, and the father of a male household head (or husband's father of female household heads) in the household file. I fill in the father's education first by using information from the resident and the non-resident files, before proceeding to fill in missing information from the household file for male household heads and spouses of female household heads.[23] The years of education range from 0 to 16, with years 1-12 indicating the

[22] An alternative would be to start with the 32,650 cases and drop only those with missing occupation while allowing for missing education. This approach would recover only 126 additional cases, so I use the stricter definition.

[23] There are some misreporting issues in the education of the father. In practise, the education of the household head's father can be recorded from three sources: (1) the resident file (if the father is co-resident); (2) from the household head's response to the question on his father's education; and (3) from his wife's response to the husband's father's education question in the women's survey file. Asher et al. (2018) show that the average correlation between father's education measured across these three sources in IHDS-II is

12 years of primary and secondary schooling, 13-15 representing years of college education,[24] and 16 signifying more than 15 years of education.[25] Figure 1 plots the average value of the years of education acquired by different birth cohorts for the OBCs and the general category separately. Even though I do not use the full sample to analyse education IGM, it is reassuring to see (roughly) parallel pre-treatment trends for this variable.

I also construct a binary variable that takes the value one if a son acquires higher years of education as compared to their father, and zero otherwise. Figure 2 depicts the percentage of individuals acquiring education greater than their father's education for different birth cohorts and caste groups. The trend for OBC and general category sons seems broadly parallel before 1971.

3.4.5.2 Education Categories

I focus on the subsamples at the extreme end of the education distribution and construct two binary variables that represent the highest level of education attained by an individual: (1) illiterate (0 years of schooling), and (2) at least high school (12 or more years of education) I use this categorization to set the data based on father's education category.

To study drop out rates of sons, I create two binary variables from the years of education variable: (1) illiterate, and (2) completed at least high school.

3.4.5.3 Occupation Categories

I ascertain the main work activity (i.e. the work activity in which an individual spends most of his labor hours) and the associated occupation codes of that activity for fathers and sons using the method delineated in appendix:occupations. Occupations in IHDS-II are classified at the two-digit level following the 1968 National Classification of Occupations (NCO) codes. NCO codes are aligned with the International Standard Classification of Occupations (ISCO) with appropriate adjustments suitable for the Indian economy. Table C1 lists the two-digit NCO-1968 codes.

0.9, but the misreporting errors are uncorrelated with household characteristics.

[24] Most undergraduate college education in India takes the form of three-year programs, but some fields like engineering have a four-year undergraduate program.

[25] Attained while pursuing a master's degree, for example.

Of course, the two-digit occupation codes in the IHDS-II will only be available for those who are employed in the labor market. Some individuals might be unemployed or not participate in the labor force, for example, due to retirement, health reasons, engagement in household work, etc. As the IHDS is a panel data set, I am able to recover the occupation codes for some individuals who were employed in 2004-05 (i.e. when the first round was conducted) but not in 2011-12 (i.e. when the second round was conducted). Ultimately, for the analysis of occupational mobility, I drop observations in which the two-digit occupation code is still not available for the son and/or the father.

Next, I aggregate the two-digit occupation codes into four categories (professional, skilled, unskilled, and farmer) by combining similar occupations. I use the classification schemes following Azam (2013) and present the two-digit occupation codes of these categories in Table C2. Figure 2 shows the percentage of sons employed in a professional or skilled occupation by caste and birth year.

3.4.5.4 Independent Variables

The two most important independent variables in my analysis are caste group affiliation and birth cohort. The caste group affiliation variable, *OBC*, is a binary variable that takes the value one if an individual reports belonging to the OBC category at the time of the survey, and zero if he belongs to the General category. If a father's caste is not explicitly recorded in the data, I assume it is the same as the son's.[26] *Young* is a binary variable that takes the value one if an individual was 21 or younger in 1993 (i.e. the time when the job quotas for OBCs were introduced), and zero otherwise. I also include the following socioeconomic variables as covariates: age of son, square of age of son, relationship of son to household head, marital status of the son, indicator for whether the household is below poverty line, religion, and place of residence (urban/rural).

[26]The OBC category is not static; the list of *jatis* classified as OBC can change over time as governments revise the official lists. However, I only observe the OBC status at the time of the survey.

3.4.6 Summary Statistics

oTable 3.2 details the mean and standard deviation of various education, occupation, father's education, father's occupation, and other socio-economic characteristics of sons belonging to my sample consisting of general category and OBC individuals who are classified as "old" (Columns (1) and (2) respectively. Column (3) tests whether the population mean of characteristics of old general-category sons is equal to those of old OBC sons. This is done to check if the levels of the outcome variables or other characteristics of the treatment and control group are similar before treatment. As can be seen from Column (3), almost all the differences are significantly different from zero. Even though I did not do a formal one-sided test, the average education and occupation outcomes of general category members are better than the OBCs. These differences do not pose a direct threat to identification in the difference-in-differences strategy that I use later. The key identification assumption is not identical baseline characteristics; the identification assumption is that in the absence of the policy, the average outcomes of OBCs would follow parallel trends to those in the general category, conditional on the included covariates.

3.5 Methodology

I partition this section into two subsections. First, I discuss the difference-in-differences framework that I use to ascertain the impact of job quotas on inter-generational mobility. The second subsection discusses some concerns in the empirical strategy.

3.5.1 Impact of Job Quotas: Difference-in-Differences Framework

I now turn to the key question of this paper: how do the education and occupation choices of sons relative to their fathers change as a result of the job quotas? I exploit the quasi-experimental nature of the implementation of the Mandal Commission reforms and use a difference-in-differences approach to compare the variation in the outcome variables of individuals affected by the program, relative to those who are not. The following two sources of variation identify an individual's exposure to the Mandal Commission reforms: his (1) caste group affiliation and (2) birth cohort. Recall that in section 3.4.3, I had specified that the treatment group is the OBCs and the control group is the General category. The treatment cohort is those who were 21 or younger at the time the job quotas were implemented. Thus, only young OBCs could be affected by the policy change, while the rest of the individuals should be unaffected.[27]

I employ various specifications in a difference-in-differences framework to see how absolute mobility changes as a result of the job quotas. I set the data based on father's education or occupation and then estimate the following regression:[28]

$$y_{iga} = \beta(obc_g \times young_a) + \gamma_{sg} + \theta_{sa} + X'_{iga}\boldsymbol{\delta} + \varepsilon_{iga} \tag{3.1}$$

where y_{icgt} is the outcome of interest of son i in caste group g and of age a; obc_g is a dummy that takes the value one if caste group is OBC and zero otherwise; $young_a$ is a dummy that takes the value one if age is less than 22 in 1993 and zero otherwise; γ_{sg} are state-caste fixed effects to take into account the state-level quotas for OBCs for employment in state-government jobs; θ_{sa} are state-cohort fixed effects to take into account time-varying state-level policies that could affect education or occupation choice decisions; X'_{iga} are covariates (age of son, square of age of son, relationship to household head, marital status, amount of land owned, religion, and place of residence (urban/rural), and ϵ_{iga} is the error term.

Standard errors are clustered at the state level. Given the small number of treated clusters (one treated state per policy evaluation), conventional inference (based on p-values from asymptotic distributions) may fail, thus, I use wild bootstrap standard errors, clustered at the state level (MacKinnon et al., 2023).

The coefficient of interest is β, and is identified under the assumption that the average outcomes of OBC sons (in a sample restricted by some indicator of the father's education or occupation) and their counterparts in the general category would follow parallel paths over birth cohorts in the absence of the job quotas (i.e. other factors that can cause these variables to change over time (e.g. trade liberalization) would affect different castes belonging to the

[27] Assuming no indirect effects to young General category individuals.

[28] I set my data in all the regressions except the one in which I consider the binary variable that a son has higher education than his father as the dependent variable.

same age cohort in the same manner).

I also estimate the event study counterpart of Equation (3.1), which allows me to visually (and flexibly) examine the differential patterns in the outcomes between the treated and controls, relative to the time of policy implementation. To increase statistical power, I collapse birth cohorts into 5-year age groups, denoted by $\hat{a}$. For instance, individuals aged 18–22 in 1993 constitute the first treated age group ($\hat{a} = 1$), while those aged 23–27 serve as the baseline group ($\hat{a} = -1$).

$$y_{iga} = \sum_{z=-4, z\neq -1}^{2} \beta_z (obc_g \times [1\{\hat{a} = z\}]) + \gamma_{sg} + \theta_{sa} + X'_{iga}\boldsymbol{\delta} + \varepsilon_{iga} \tag{3.2}$$

I now discuss the exact y_{iga} variable used and the method of sub-setting the data for studying education and occupation mobility. As will be evident shortly, restricting the sample based on some indicator of father's education or occupation allows us to interpret β as changes in absolute mobility.[29]

3.5.1.1 Effect on Education IGM

I use three different methods to study the effect of job quotas on absolute education IGM. First, I look at the effect of the policy on the probability that a son acquires higher years of education as compared to their father. I do this analysis using my entire sample. However, the β coefficient in this case gives us the impact of the policy on average and could mask important heterogeneity. For example, if we observe β to be positive (which would be an indication of upward absolute education mobility), could it be possible that sons of highly-educated fathers are acquiring more education than their (highly educated fathers) and sons of lower-educated fathers are not? To study this, I use subsamples constructed using the highest education level attained by fathers.

Second, I use the two subsamples based on father's education categories defined in Section 3.4.5.2 (the first one consisting of illiterate fathers and the second one consisting of fathers who have completed at least high school) and for each of the 2 subsamples, I estimate

[29] An alternative way could be to interact obc_g, $young_a$, and an indicator for father's education or occupation category. However, I follow the prior literature, such as Alesina et al. (2021) and Asher et al. (2018), that seems to prefer the subsectioning method.

Equation (3.1) separately with y_{icgt} being the son's years of education variable as defined in Section 3.4.5.1.

Third, I analyse inter-generational education transitions. I focus on the two most important transitions. First, I look at the effect of the policy on the probability that the son is literate conditional on his father being illiterate. A positive β in this case would signal absolute upward education mobility. Second, I look at the effect of the policy on the probability that the son does not complete high school conditional on his father being at least a high school graduate. A positive β in this case would signal absolute downward education mobility.

3.5.1.2 Effect on Occupation IGM

I now turn to inter-generational occupation mobility. I use occupation transition probabilities to study occupation IGM. The conditional probability of an occupation transition from the father's generation to the son's generation is obtained in a manner similar to the education-transition probabilities. Thus, in Equation (3.1), y_{iga} now represents the occupation category instead of the education category and the sample is partitioned into subsamples using father's occupation categories. Occupation categories are defined in Section 3.4.5.3, but I combine (1) professional and skilled and (2) unskilled and farmer occupations together instead of producing the effect of job quotas using a 4×4 occupation transition matrix. I do this because it is not immediately obvious how farming and unskilled jobs rank against each other (for education, more years of education would be considered better, and thus there is no ambiguity in the interpretation of a full educational mobility matrix).

To study absolute upward occupation mobility, I use the probability that a son who is born to a father employed in an unskilled or farming occupation is employed in a professional or skilled occupation. Similarly, to study absolute downward occupation mobility, I use the probability that a son who is born to a father employed in an professional or skilled occupation is employed as a farmer or in an unskilled job.

3.5.2 Possible Concerns in the Difference-in-Differences Analysis

In an ideal world, we would use well designed experiments to uncover causal effects of policy interventions. However, in the real world, randomized experiments are costly or just infeasible in many substantive domains of interest, so economists use empirical design-based tools that widen the range of possible variation that can be used to uncover causal impacts. Difference-in-differences is "probably the most widely applicable design-based estimator" (Angrist and Pischke, 2010), but its credibility rests on some identifying assumptions. I now list the main identification concerns in the use of a difference-in-differences estimator to explain the effect of job quotas on the IGM of the OBCs.

First, it is possible that the conditional parallel trends assumption (which requires that, after controlling for covariates, the control group serves as a valid counterfactual for the treated group) does not hold. Specifically, the outcomes for OBC children may have evolved differently from those of general category children even in the absence of the quota policy. To assess this, I visually inspect the pre-treatment coefficients from Equation 3.2.

The second concern is the assumption that there are no spillover effects of the policy change on caste groups other than the OBCs. A widely-held belief regarding reservations in India is that they harm the non-targeted social groups i.e. the forward castes by effectively decreasing the percentage of seats available for them in government jobs and higher educational institutions. Indeed, any proposed increase in reserved seats is met by large-scale protests and strikes. For e.g., the implementation of the Mandal Commission Reforms, the specific affirmative action policy that I consider, was met by violent protests in the 1990s in which some upper-caste students even self-immolated themselves. Thus, the upper-caste students might feel discouraged by the resultant increase in competition and work less hard,[30] or alternatively, they might increase their effort to be able to secure these seats[31]. When spillover effects are suspected, identifying a suitable control group in a non-experimental research design is, therefore, further complicated by having to identify a subject group (a subset of the ineligible groups of the treatment policy) for which spillover effects are pos-

[30] If the true effect of the policy is positive, then the estimates would be positively biased in this case as the outcomes of the control group individuals are more likely to be worse than in the absence of spillover effects.

[31] In this case, the estimates would be negatively biased.

sible and a control group that won't be affected by the treatment at all. The other two caste groups in India, the SCs and the STs, should not be affected by the implementation of the job quotas for the OBCs as the percentage of seats reserved for them did not change. However, as mentioned in Section 3.4.3, they also experienced a policy change that affected their access to public resources at the same time as the implementation of the job quotas for the OBCs, and thus, can't be suitable control groups. The suspected spillover effects are offset to some degree by political promises that the total number of government jobs would be increased so that the absolute number of jobs available to the upper-caste members is unchanged (Khanna, 2020), but the difference-in-differences model will prove to be inadequate if the spillover effects still remain large. As a more recent example, the government implemented a 10% quota in educational institutions for economically weaker sections of the general category in 2019. To ensure that the pre-2019 reservation remains unaffected, the government plans to increase the total seats in educational institutions by as much as 25%.

The third assumption is that there are no anticipatory effects of the policy. Given the long and tumultuous journey of the OBC quotas, it is unlikely that households actually expected the quotas to be actually implemented. As mentioned in Section 3.3.3, the initial impetus for OBC quotas came in the 1950s, but no action was taken by the government for decades. Even when the government first announced their intention to act upon the recommendations of the Mandal commission, for a couple of years the fate of the policy pretty much hung in the balance due to the widespread anti-reservation protests. Finally, whether the quotas would be imposed or not was ultimately decided by the Supreme Court. Indian courts are notorious for cases being dragged on for several years, even decades, and thus, it was certainly a surprise that the verdict was announced within one year. Hence, I do not expect that the prospect of the policy being adopted changed between 1989 and 1993.[32]

[32] A robustness check could be to assume that cohorts younger than 22 in 1989 rather than 1993 would be affected by the policy. If the results from this treatment year do not change much in magnitude as compared to the results using 1993 as the treatment year, then it would suggest that the policy had no anticipatory effects.

3.6 Results

I divide this section into two subsections. I first present results on the effect of the implementation of the Mandal Commission Reforms in 1993 on absolute education mobility. The second subsection contains results on occupation related outcomes.

3.6.1 Education Mobility

Table 3.3 presents the impact of the OBC job quotas on education mobility related measures. Table 3.3 is read as follows. Each column of the table represents a different OLS regression defined by the dependent variable and the sample used. All regressions include state–age fixed effects, state–caste fixed effects, and the set of covariates described in Section 3.4.5.4. For brevity, I omit the coefficients on the fixed effects and covariates, and report only the coefficient of interest on the interaction term $obc_g \times young_a$, β, along with its standard error, the control group mean of the dependent variable, and the number of observations.

Column (1) estimates Equation 3.1 using as the dependent variable an indicator for whether the son attained more education than his father. The estimated coefficient is positive and statistically significant, implying that the quotas increased the probability of upward educational mobility by 4.2 percentage points (a 6.1 percent increase relative to the baseline).

However, as mentioned in Section 3.5.1.1, this number could mask important heterogeneity. Thus, Columns (2) and (3) restrict the sample to sons of illiterate fathers, to assess mobility from the very bottom of the education distribution. In Column (2), the dependent variable is the son's years of schooling, while in Column (3) it is an indicator for whether the son is literate. The estimated effect on years of schooling is positive and sizable: sons of illiterate fathers exposed to the quotas attained 0.62 additional years of schooling on average, which represents a roughly 10 percent increase relative to the pre-policy mean. In contrast, the effect on literacy is positive but not statistically significant

Columns (4) and (5) examine the subsample of sons whose fathers had completed at least high school. Column (4) uses the son's years of education as the dependent variable, while Column (5) uses an indicator for whether the son fails to complete high school. In this relatively advantaged group, the estimated effects are small and statistically indistinguishable from zero. The point estimates suggest an increase of 0.22 years of schooling and a 7.9 percentage point decline in the probability of not completing high school, but neither estimate is precise enough to rule out modest positive or negative effects.

Panels (a)–(e) of Figure 3.3 present the event-study coefficients from Equation 3.2, where I group sons into five-year age cohorts (measured in 1993, the year of quota implementation) and omit the 23–27 cohort as the reference category. Across all specifications, the event-study plots show no evidence of differential pre-trends between OBCs and others, providing support for the conditional parallel trends assumption. The estimated effects are largest in magnitude for the youngest cohorts in my sample (those aged 8–12 in 1993), who had the greatest scope to adjust their schooling decisions in response to the policy, relative to older cohorts whose education trajectories were already more fixed.

3.6.2 Occupation Mobility

Table 3.4 contains results of the effect of the policy on inter-generational occupation transitions. In contrast to the gains in educational mobility, the results in Table 3.4 indicate a slight decline in upward occupational mobility. Column (1) of Table 3.4 restricts the sample to sons born to fathers engaged in farming or other unskilled occupations. The dependent variable used in the estimation of Equation 3.1 is a binary variable that takes the value one if the son's occupation is professional or skilled, and zero otherwise. For OBC sons born to fathers in unskilled or farming occupations, the probability of transitioning to professional or skilled jobs decreased by 4.9 percentage points, indicating downward absolute occupation mobility. Panel (a) of Figure 3.4 plots the corresponding event study, showing that this decline remains relatively stable across cohorts of OBC youth. At the same time, there is a 4.79 percentage point increase in the probability of employment in government jobs. However, not all government jobs are "skilled"; for example, many service-sector positions such as janitors fall into this category.

In contrast to the gains in educational mobility, Table 3.4 shows a decline in upward occupational mobility. Panel (a) of Figure (3.4 plots the correspnding event study and shows the decline is more or less stable for all cohorts of OBC youth. There is an increase in the percentage employed in govt. jobs by 4.79 percentage points.

Column (3) restricts the sample to sons of fathers employed in skilled or professional occupations. Here, the dependent variable is a binary indicator equal to one if the son's occupation is skilled or professional, and zero otherwise. For this group, the probability increased by 1.44 percentage points, though the estimate is statistically insignificant. In Column (4), I use government employment as the dependent variable, and find that the probability of obtaining a government job rises by 7.98 percentage points. Panels (b) and (d) of Figure 3.4 show stronger increases in government employment among younger cohorts, both for sons of fathers in unskilled or farming occupations and for sons of fathers in skilled or professional occupations.

Taken together, these results suggest at least one clear piece of evidence of absolute downward occupational mobility following the introduction of the quotas, even though access to government employment increased. One potential explanation is that the quotas exacerbated negative stereotypes among private employers toward OBC workers. As a result, while OBCs experienced greater opportunities in the public sector, their opportunities in the private sector may have declined.

Finally, in Table 3.5, I examine the effects of the policy for son–father pairs with missing occupational information for at least one generation. I find no statistically significant effect on the probability that the son's occupation is missing when the father's occupation is unskilled (Column 1) or skilled (Column 2). Likewise, there is no significant effect on the probability that the son's occupation is skilled, in government, or missing (Columns (3), (4), and (5), respectively) when the father's occupation is missing.

3.7 Conclusion

Centuries-old caste-based discriminatory practises and attitudes have given rise to heterogeneity in the existing socio-economic status in India, with only some who are able to climb the socio-economic ladder, while the rest either remain at the same position in society or find their position declining (Ramaiah, 1992). Affirmative action policies, enacted as reservation quotas in jobs, higher education institutes, and legislatures, try to correct this imbalance. Using a sample of father–son pairs from IHDS-II and a difference-in-differences strategy, I find that the introduction of job quotas in the public sector for the OBCs led to an increase in their absolute upward educational mobility, but a slight decline in their absolute occupational mobility. Public-sector employment among OBCs rose, and across all outcomes considered, the effects are strongest for the youngest OBC cohorts in my sample.

These findings highlight a complex dynamic: while the policy succeeded in raising educational attainment, likely due to improved expectations about future opportunities, it did not translate into occupational mobility. Even with increased access to public-sector jobs, OBCs may have faced limited opportunities or heightened discrimination in private labor markets or switched to low-skilled public sector jobs. Thus, public-sector quotas alone may be insufficient to close long-term occupational gaps without complimentary efforts in the private sector.

3.8 Figures and Tables

Figures

Figure 3.1: Years of Education

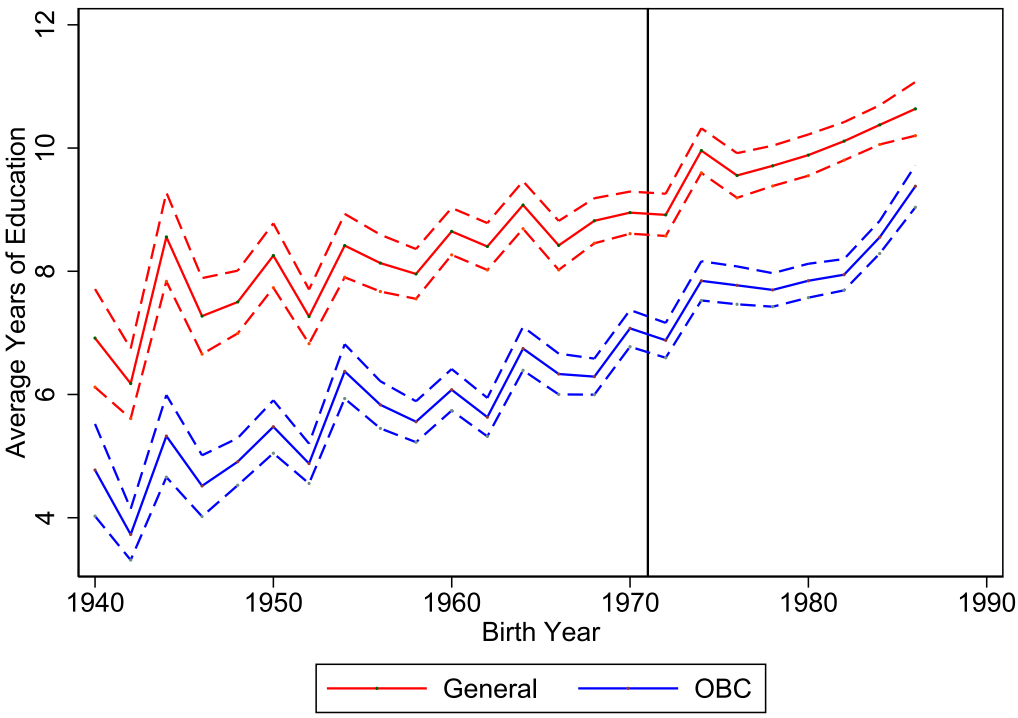

Note: This figure plots the average value of the years of education acquired by different birth cohorts. See Section 3.4.5.1 for variable definition. The solid red line denotes average education for the general caste category, and the blue line denotes the average education for the OBCs. Dashed lines represent the 95% confidence intervals. The solid black line at birth year= 1971 indicates the earliest birth cohorts treated by the implementation of the job quotas.

Figure 3.2: Percentage of Individuals in Professional/Skilled Jobs over Time

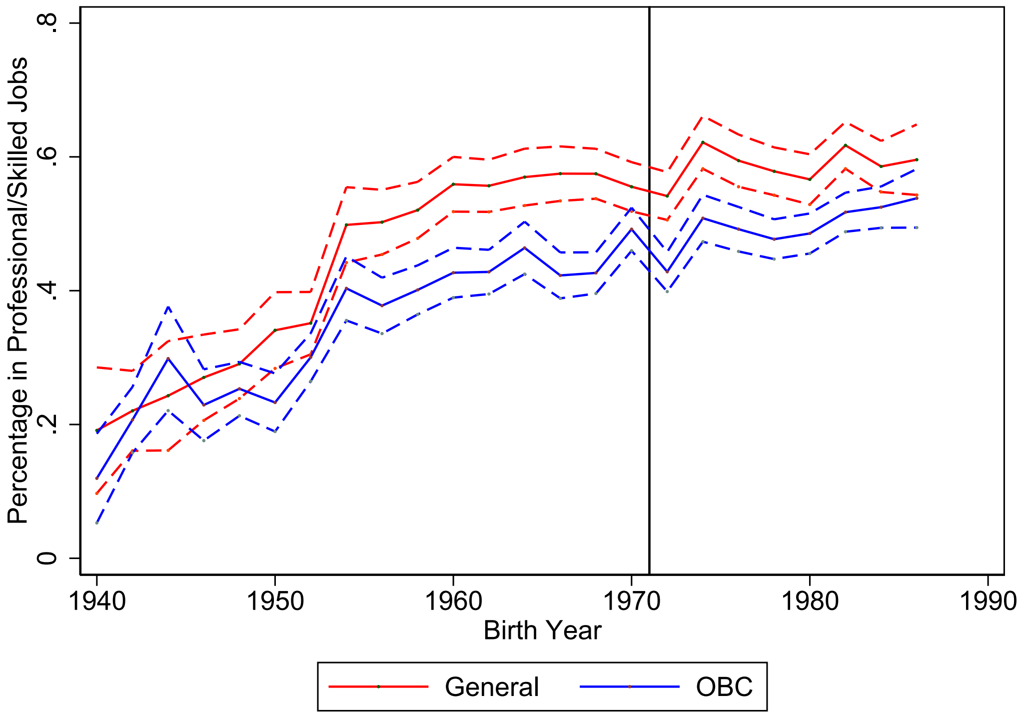

Note: This figure plots the percentage of individuals employed in professional or skilled occupations for different birth cohorts. See Section 3.4.5.3 for variable definition and construction. The solid red line denotes this percentage for the general caste category, and the blue line denotes this percentage for the OBCs. Dashed lines represent the 95% confidence intervals. The solid black line at birth year= 1971 indicates the earliest birth cohorts treated by the implementation of the job quotas.

Figure 3.3: Effects of 1993 job quotas on education mobility

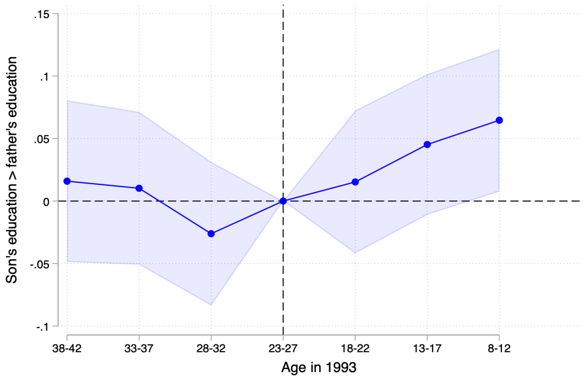

(a) All

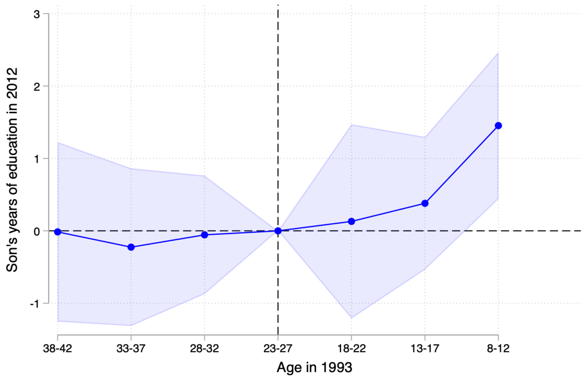

(b) Sons with illiterate fathers

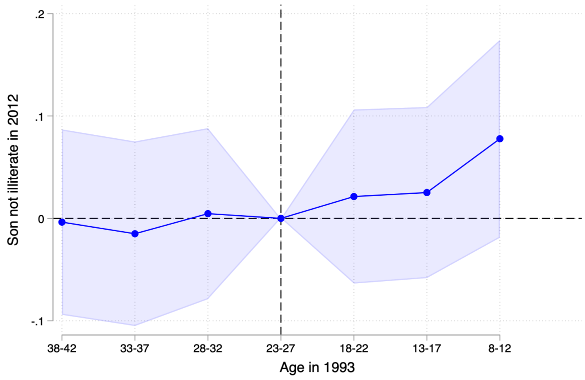

(c) Sons with illiterate fathers

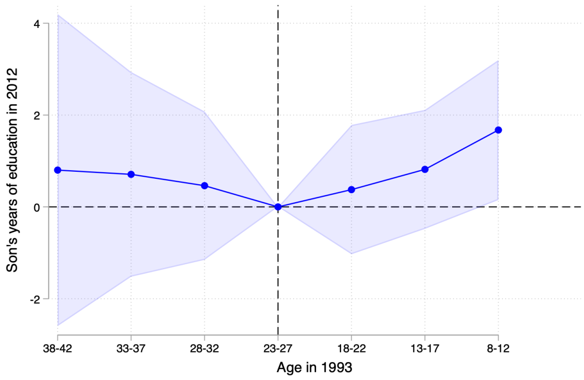

(d) Sons with high school graduate fathers

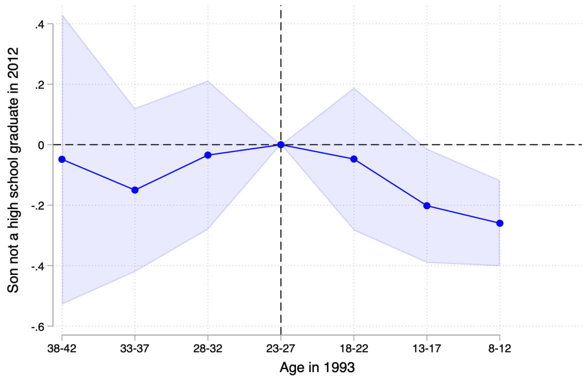

(e) Sons with high school graduate fathers

Note: This figure reports event-study estimates of the effect of the 1993 OBC job quotas on education-related outcomes. Panel (a) shows the probability that a son's education exceeds his father's (full sample). Panels (b)–(c) show the son's years of education and the probability that the son is illiterate, restricting to sons with illiterate fathers. Panels (d)–(e) show the son's years of education and the probability that the son is not a high-school graduate, restricting to sons with high-school-graduate fathers. Each panel plots the estimated β_z coefficients from equation (3.2) with 95% confidence intervals. All regressions include fixed effects for state-caste, state-birth years, plus controls (religion, son's age, urban/rural status, etc.).

Figure 3.4: Effects of 1993 job quotas on occupation mobility

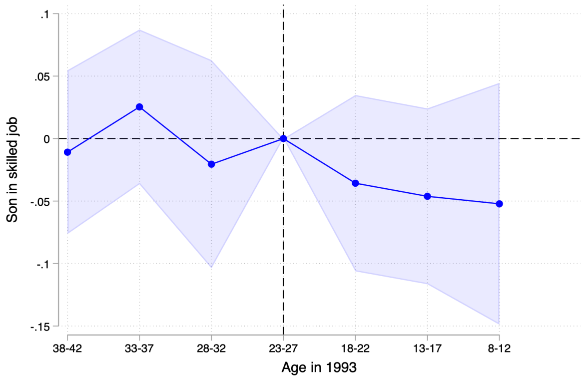

(a) Sons with fathers in unskilled occupations

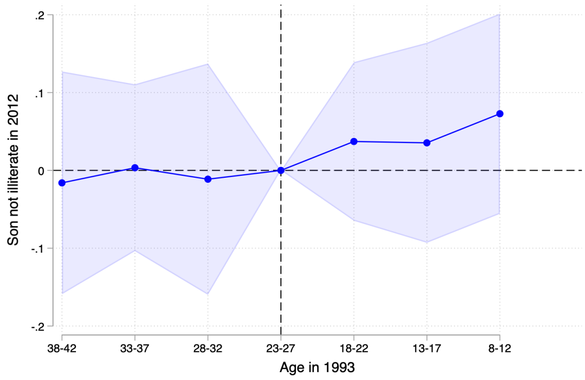

(b) Sons with fathers in unskilled occupations

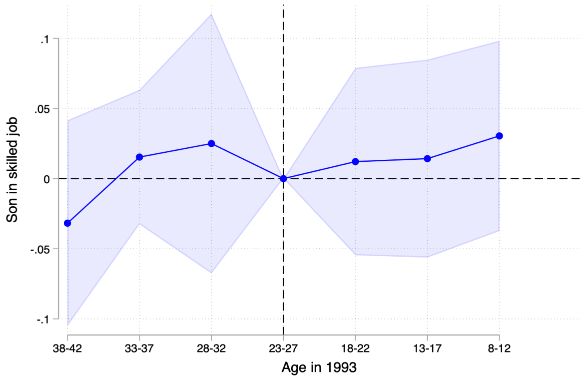

(c) Sons with fathers in skilled occupations

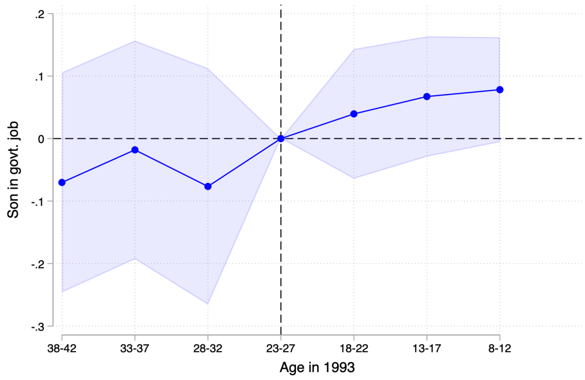

(d) Sons with fathers in skilled occupations

Note: This figure reports event-study estimates of the effect of the 1993 OBC job quotas on occupation-related outcomes. Panels (a)–(b) show the probability that the son is in a skilled and govt. job, respectively, restricting to sons with illiterate fathers. Panels (c)–(d) show the son is in a skilled and govt. job, respectively, restricting to sons with high-school-graduate fathers. Each panel plots the estimated β_z coefficients from equation (3.2) with 95% confidence intervals. All regressions include fixed effects for state-caste, state-birth years, plus controls (religion, son's age, urban/rural status, etc.).

Table 3.1: Sample Restrictions

Panel A: Initial sample restrictions

Restriction	Resident son with residen-t/dead father	Resident son with nonresi-dent father	Nonresident son with resident father	Total
Keeping only males	99,071	2,991	11,376	113,438
Keep OBC/general category	68,167	2,237	8,377	78,781
Keep those whose age is not missing and are younger by at least 10 years than their father	68,150	2,233	8,256	78,639
Keep birth cohorts 1951–1985	28,361	119	4,170	32,650

Panel B: Restrictions to get son-father pairs for studying education mobility

Sample restrictions	General old	OBC old	General young	OBC young	Total
No restriction	6,934	9,086	6,811	9,819	32,650
Resident son with resident/dead father	6,328	8,305	5,661	8,067	28,361
Resident son with nonresident father	6	12	38	63	119
Nonresident son with resident father	600	769	1,112	1,689	4,170
Keep those not currently in school	6,934	9,086	6,772	9,786	32,578
Son education is not missing	6,932	9,076	6,757	9,760	32,525
Father education is not missing	5,919	7,882	5,137	7,529	26,467
Father's education recovered because son is a HH head (even if father is dead)	5,134	7,115	1,983	3,657	17,889
Son-father pairs for studying education mobility	5,919	7,882	5,137	7,529	26,467

Panel C: Further restrictions to get son-father pairs for studying occupation mobility

Sample restrictions	General old	OBC old	General young	OBC young	Total
Son's occupation is not missing	5,538	7,449	4,838	7,186	25,011
Father's occupation is not missing	5,191	7,169	4,046	6,356	22,762
Father's occupation recovered because son is a HH head (even if father is dead)	4,947	6,954	2,083	3,851	17,835
Son-father pairs for studying occupation mobility	5,191	7,169	4,046	6,356	22,762

Note: Table reports the number of son-father pairs available in the IHDS-II survey for studying education and occupation mobility.

Table 3.2: Summary Statistics

	(1) Old General mean/sd	(2) Old OBC mean/sd	(3) Diff b/t
Son's age in 1993	31.56	31.53	0.03
	(5.69)	(5.76)	0.29
Son's education (in years)	8.45	6.09	2.35***
	(4.89)	(4.72)	27.81
Father's education (in years)	4.35	2.38	1.97***
	(4.66)	(3.60)	26.47
Percentage of illiterate sons	0.14	0.26	−0.11***
	(0.35)	(0.44)	−16.71
Percentage of illiterate fathers	0.42	0.61	−0.18***
	(0.49)	(0.49)	−21.36
Percentage of high school graduate sons	0.28	0.13	0.15***
	(0.45)	(0.34)	20.96
Percentage of high school graduate fathers	0.09	0.03	0.06***
	(0.28)	(0.16)	13.95
Percentage of sons in professional jobs	0.16	0.08	0.08***
	(0.36)	(0.27)	12.66
Percentage of fathers in professional jobs	0.09	0.04	0.05***
	(0.28)	(0.19)	11.28
Percentage of sons in skilled jobs	0.38	0.34	0.04***
	(0.48)	(0.47)	4.65
Percentage of fathers in skilled jobs	0.27	0.28	0.00
	(0.45)	(0.45)	−0.35
Percentage of sons in unskilled jobs	0.10	0.17	−0.06***
	(0.30)	(0.37)	−10.59
Percentage of fathers that are in unskilled jobs	0.16	0.23	−0.07***
	(0.37)	(0.42)	−9.40
Percentage of sons that are farmers	0.36	0.41	−0.05***
	(0.48)	(0.49)	−5.86
Percentage of fathers that are farmers	0.48	0.46	0.02*
	(0.50)	(0.50)	2.07
Percentage of sons in govt. jobs	0.37	0.22	0.16***
	(0.48)	(0.41)	14.28
Percentage of married sons	0.96	0.96	0.00
	(0.19)	(0.20)	0.44
Percentage of sons who are Hindu	0.76	0.84	−0.08***
	(0.43)	(0.37)	−11.47
Percentage of sons below poveerty line	0.08	0.13	−0.05***
	(0.27)	(0.33)	−9.44
Percentage of sons in urban areas	0.46	0.37	0.10***
	(0.50)	(0.48)	11.32
N	5689	7548	13237

Note: Table reports the baseline summary statistics by caste.

Table 3.3: Effect of 1993 job quotas on education related variables (DiD estimates)

	More years	Years of education	Not illiterate	Years of education	Not high school graduate
OBCxYoung	0.0417**	0.6157*	0.0317	0.2210	−0.0794
	(0.0212)	(0.3486)	(0.0322)	(0.4745)	(0.0629)
Dep var mean	0.6881	5.6930	0.7071	13.6820	0.1724
Observations	25,188	11,734	11,734	1,582	1,582
Sample	All	Iliterate fathers		High school graduate fathers	

Note: This table shows the point estimate and the standard error of the effect of the 1993 job quotas for OBCs on education related variables. Each column is a separate regression and is estimated using the difference in differences equation (3.1).

Table 3.4: Effect of 1993 job quotas on occupation related variables (DiD estimates)

	Son in skilled/profes-sional	Son in govt.	Son in skilled/professional	Son in govt.
OBCxYoung	−0.0490**	0.0479**	0.0144	0.0798**
	(0.0186)	(0.0237)	(0.0270)	(0.0361)
Dep var mean	0.3555	0.3461	0.8305	0.3974
Observations	14,259	8,500	7,571	4,649
Sample	Unskilled/farmer fathers		Skilled/professional fathers	

Note: This table shows the point estimate and the standard error of the effect of the 1993 job quotas for OBCs on occupation related variables. Each column is a separate regression and is estimated using the difference in differences equation (3.1).

Table 3.5: Effect of 1993 job quotas on occupation related variables (DiD estimates)

	Son's occupation missing	Son's occupation missing	Son in skilled/pro-fessional	Son in govt.	Son's occupation missing
OBCxYoung	0.0048	0.0036	−0.0804	−0.0736	0.0475
	(0.0137)	(0.0099)	(0.0609)	(0.0962)	(0.0440)
Dep var mean	0.0480	0.0854	0.6974	0.4122	0.1034
Observations	14,795	8,107	1,708	1,013	1,919
Sample	Unskilled/ farmer fathers	Skilled/ professional fathers	Fathers with missing occupation		

Note: This table shows the point estimate and the standard error of the effect of the 1993 job quotas for OBCs on education related variables. Each column is a separate regression and is estimated using the difference in differences equation (3.1).

Bibliography

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.

_ **and Javier Gardeazabal**, "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, 2003, *93* (1).

Abel, Martin, Eliana Carranza, and Maria Elena Ortega, "Can Temporary Wage Incentives Increase Formal Employment? Experimental Evidence from Mexico," *IZA Discussion Paper No. 15740*, 2022.

Ahn, Namkee, "Measuring the Value of Children by Sex and Age Using a Dynamic Programming Model," *The Review of Economic Studies*, 1995, *62* (3), 361–379.

Aizawa, Naoki, Soojin Kim, and Serena Rhee, "Labor Market Screening and the Design of Social Insurance: An Equilibrium Analysis of the Labor Market for the Disabled," *Review of Economic Studies*, 2025, *92* (1), 1–39.

Alesina, Alberto, Sebastian Hohmann, Stelios Michalopoulos, and Elias Papaioannou, "Intergenerational Mobility in Africa," *Econometrica*, 2021, *89* (1), 1–35.

Ananat, Elizabeth Oltmans, Jonathan Gruber, Phillip Levine, and Douglas Staiger, "Abortion and Selection," *The Review of Economics and Statistics*, 2009, *91* (1), 124–136.

Anderson, Siwan and Debraj Ray, "Missing Women: Age and Disease," *The Review of Economic Studies*, 2010, *77* (4), 1262–1300.

Angrist, Joshua and Jörn-Steffen Pischke, "The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics," *Journal of Economic Perspectives*, 2010, *24* (2), 3–30.

Anukriti, S, "Financial Incentives and the Fertility-Sex Ratio Trade-Off," *American Economic Journal: Applied Economics*, 2018, *10* (2), 27–57.

Arnold, Fred, "Sex Preference and Its demographic and health implications," *International Family Planning Perspectives*, 1992, pp. 93–101.

_ , Minja Kim Choe, and Tarun Roy, "Son Preference, the Family-Building Process and Child Mortality in India," *Population Studies*, 1998, *52* (3), 301–315.

Asher, Sam, Paul Novosad, and Charlie Rafkin, "Intergenerational Mobility in India: Estimates from New Methods and Administrative Data," *World Bank Working Paper*, 2018, *10* (3), 642–662.

Autor, David H and Susan N Houseman, "Do Temporary-Help Jobs Improve Labor Market Outcomes for Low-Skilled Workers? Evidence from "Work First," *American Economic Journal: Applied Economics*, 2010, *2* (3), 96–128.

Azam, Mehtabul, "Intergenerational Occupational Mobility in India," *IZA Discussion Paper No. 7608*, 2013.

_ and Vipul Bhatt, "Like Father, Like Son? Intergenerational Educational Mobility in India," *Demography*, 2015, *52* (6), 1929–1959.

Bali, Surya, Kriti Yadav, and Yash Alok, "A Study of Quality of Care of Sterilization Services in Madhya Pradesh," *Journal of Family Medicine and Primary Care*, 2020, *9* (12), 6005–6011.

Banerjee, Abhijit and Thomas Piketty, "Top Indian Incomes, 1922–2000," *The World Bank Economic Review*, 2005, *19* (1), 1–20.

Bardhan, Pranab, Dilip Mookherjee, and Monica Parra Torrado, "Impact of Political Reservations in West Bengal Local Governments on Anti-Poverty Targeting," *Journal of Globalization and Development*, 2010, *1* (1).

Barnow, Burt and Jeffrey Smith, "Employment and Training Programs," in "Economics of Means-Tested Transfer Programs in the United States, Volume 2," University of Chicago Press, 2015, pp. 127–234.

Bartik, Timothy, *Jobs for the Poor: Can Labor Demand Policies Help?*, New York: Russell Sage Foundation, 2001.

_ **and George Erickcek**, "Simulating the Effects of the Tax Credit Program of the Michigan Economic Growth Authority on Job Creation and Fiscal Benefits," *Economic Development Quarterly*, 2014, *28* (4), 314–327.

Becker, Gary, "An Economic Analysis of Fertility," in "Demographic and Economic Change in Developed Countries," Columbia University Press, 1960, pp. 209–240.

_ **and Gregg Lewis**, "On the Interaction Between the Quantity and Quality of Children," *Journal of Political Economy*, 1973, *81* (2), 279–288.

Bell, Stephen and Larry Orr, "Is Subsidized Employment Cost Effective For Welfare Recipients? Experimental Evidence from Seven State Demonstrations," *Journal of Human Resources*, 1994, pp. 42–61.

Bhalotra, Sonia, Abhishek Chakravarty, and Selim Gulesci, "The Price of Gold: Dowry and Death in India," *Journal of Development Economics*, 2020, *143*, 102413.

_ **and Tom Cochrane**, "Where Have All the Young Girls Gone? Identification of Sex Selection in India," *IZA Discussion Paper No. 5381*, 2010.

Bhartiya Janta Party, "Election Manifesto 2014," `https://www.bjp.org/images/pdf_2014/full_manifesto_english_07.04.2014.pdf` 2014. Accessed May 27, 2025.

Bishop, John, "Employment in Construction and Distribution Industries: The Impact of the New Jobs Tax Credit," in "Studies in Labor Markets," University of Chicago Press, 1981, pp. 209–246.

_ **and Mark Montgomery**, "Does the Targeted Jobs Tax Credit Create Jobs at Subsidized Firms?," *Industrial Relations*, 1993, *32* (3), 289–306.

Biswas, Nabaneeta, Christopher Cornwell, and Laura Zimmermann, "The Power of Lakshmi: Monetary Incentives for Raising a Girl," *Journal of Human Resources*, 2023.

Bloch, Francis and Vijayendra Rao, "Terror as a Bargaining Instrument: A Case Study of Dowry Violence in Rural India," *American Economic Review*, 2002, *92* (4), 1029–1043.

Boockmann, Bernhard, Thomas Zwick, Andreas Ammermüller, and Michael Maier, "Do Hiring Subsidies Reduce Unemployment among Older Workers? Evidence from Natural Experiments," *Journal of the European Economic Association*, 2012, *10* (4), 735–764.

Borker, Girija, Jan Eeckhout, Nancy Luke, Shantidani Minz, Kaivan Munshi, and Soumya Swaminathan, "Wealth, Marriage, and Sex Selection," *Cowles Foundation Discussion Paper No 2814*, 2024.

Bureau of Labor Statistics, "Characteristics of Minimum Wage Workers, 2022," 2023.

_ , "May 2023 Monthly Labor Review," *Monthly Labor Review*, 2023.

Burtless, Gary, "Are Targeted Wage Subsidies Harmful? Evidence from a Wage Voucher Experiment," *ILR Review*, 1985, *39* (1), 105–114.

Bushway, Shawn and Justin Pickett, "Direct Incentives May Increase Employment of People with Criminal Records," *Criminology & Public Policy*, 2024.

Byker, Tanya and Italo Gutierrez, "Estimating the Impact of Sterilization under a Government Campaign Using Contaminated Treatment Data," *Economic Development and Cultural Change*, 2021, *70* (1), 159–202.

Cahuc, Pierre, Stéphane Carcillo, and Thomas Le Barbanchon, "The Effectiveness of Hiring Credits," *Review of Economic Studies*, 2019, *86* (2), 593–626.

Callaway, Brantly and Pedro Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Calonico, Sebastian, Matias Cattaneo, and Max Farrell, "On the Effect of Bias Estimation on Coverage Accuracy in Nonparametric Inference," *Journal of the American Statistical Association*, 2018, *113* (522), 767–779.

_ , _ , and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

_ , _ , Max Farrell, and Rocio Titiunik, "rdrobust: Software for Regression-Discontinuity Designs," *The Stata Journal*, 2017, *17* (2), 372–404.

Calvi, Rossella, Andrew Beauchamp, and Scott Fulford, "Terms of Engagement: Migration, Dowry, and Love in Indian Marriages," *CEPR Discussion Paper No. DP16659*, 2022.

Cappelli, Peter, "Assessing the Effect of the Work Opportunity Tax Credit," 2012.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

_ , Ciprian Domnisoru, and Lowell Taylor, "The Intergenerational Ttransmission of Human Capital: Evidence from the Golden Age of Upward Mobility," *Journal of Labor Economics*, 2022, *40* (S1), S39–S95.

_ , Jochen Kluve, and Andrea Weber, "Active Labour Market Policy Evaluations: A Meta-Analysis," *The Economic Journal*, 2010, *120* (548), F452–F477.

Carro, Jesús and Pedro Mira, "A Dynamic Model of Contraceptive Choice of Spanish Couples," *Journal of Applied Econometrics*, 2006, *21* (7), 955–980.

Cassan, Guilhem, Daniel Keniston, and Tatjana Kleineberg, "A Division of Laborers: Identity and Efficiency in India," *National Bureau of Economic Research No. w28462*, 2021.

Cattaneo, Matias, Michael Jansson, and Xinwei Ma, "Manipulation Testing Based on Density Discontinuity," *The Stata Journal*, 2018, *18* (1), 234–261.

_ , _ , and _ , "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chandavarkar, Rajnarayan, *The Origins of Industrial Capitalism in India: Business Strategies and the Working Classes in Bombay, 1900-1940*, Vol. 51, Cambridge University Press, 2003.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States," *The Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

Chiplunkar, Gaurav and Jeffrey Weaver, "Marriage Markets and the Rise of Dowry in India," *Journal of Development Economics*, 2023, *164*, 103115.

Clark, Shelley, "Son Preference and Sex Composition of Children: Evidence from India," *Demography*, 2000, *37* (1), 95–108.

Coate, Stephen and Glenn Loury, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?," *The American Economic Review*, 1993, pp. 1220–1240.

Cohen, Alma, Rajeev Dehejia, and Dmitri Romanov, "Financial Incentives and Fertility," *Review of Economics and Statistics*, 2013, *95* (1), 1–20.

Congress, U.S., "H.R.6833 - Improve and Enhance the Work Opportunity Tax Credit Act," https://www.congress.gov/bill/117th-congress/house-bill/6833 2023. Accessed: 2025-01-24.

Corwin, Emily, "A Tax Credit Was Meant to Help Marginalized Workers Get Permanent Jobs. Instead It's Subsidizing Temp Work.," Technical Report, ProPublica Report 2022.

Crépon, Bruno and Gerard J Van Den Berg, "Active Labor Market Policies," *Annual Review of Economics*, 2016, *8*, 521–546.

Cullen, Zoë and Bobak Pakzad-Hurson, "Equilibrium Effects of Pay Transparency," *Econometrica*, 2023, *91* (3), 765–802.

_ **, Will Dobbie, and Mitchell Hoffman**, "Increasing the Demand for Workers with a Criminal Record," *Quarterly Journal of Economics*, 2023, *138* (1), 103–150.

Cunha, Flavio and James Heckman, "The Technology of Skill Formation," *American Economic Review*, 2007, *97* (2), 31–47.

Dasgupta, Monica, "Explaining Asia's "Missing Women": A New Look at the Data," *Population and Development Review*, 2005, *31* (3), 529–535.

_ **, Jiang Zhenghua, Li Bohua, Xie Zhenming, Woojin Chung, and Bae Hwa-Ok**, "Why is Son Preference So Persistent in East and South Asia? A Cross-Country Study of China, India and the Republic of Korea," *The Journal of Development Studies*, 2003, *40* (2), 153–187.

de la Rupelle, Maëlys and Christelle Dumas, "Sterilizations and Women Health in India," *World Development*, 2025, *193*, 107020.

Department of Labor, "Work Opportunity Tax Credit; Request for Comments Regarding Proposed Modifications to Procedural," 2023.

_ , "WOTC Certifications by Recipient Group: State and National Details For Fiscal Year 2022," 2023.

_ , "WOTC Certifications by Recipient Group: State and National Details For Fiscal Year 2023," 2024.

Derenoncourt, Ellora, "Can you move to opportunity? Evidence from the Great Migration," *American Economic Review*, 2022, *112* (2), 369–408.

Desai, Jaikishan and Alessandro Tarozzi, "Microcredit, Family Planning Programs, and Contraceptive Behavior: Evidence from a Field Experiment in Ethiopia," *Demography*, 2011, *48* (2), 749–782.

Drèze, Jean and Mamta Murthi, "Fertility, Education, and Development: Evidence from India," *Population and Development Review*, 2001, *27* (1), 33–63.

Dyson, Tim and Mick Moore, "On Kinship Structure, Female Autonomy, and Demographic Behavior in India," *Population and Development Review*, 1983, *9* (1), 35–60.

Fenizia, Alessandra, Nicholas Li, and Luca Citino, "The (In)effectiveness of Targeted Payroll Tax Reductions," *SSRN 5097079*, 2025.

Galasso, Emanuela, Martin Ravallion, and Agustin Salvia, "Assisting the Transition From workfare to Work: A Randomized Experiment," *ILR Review*, 2004, *58* (1), 128–142.

Garin, Andrew, Dmitri Koustas, Carl McPherson, Samuel Norris, Matthew Pecenco, Evan Rose, Yotam Shem-Tov, and Jeffrey Weaver, "The Impact of Incarceration on Employment, Earnings, and Tax Filing," *Econometrica*, 2025, *93* (2), 503–538.

Gelber, Alexander, Adam Isen, and Judd Kessler, "The Effects of Youth Employment: Evidence from New York City Lotteries," *Quarterly Journal of Economics*, 2016, *131* (1), 423–460.

González, Libertad and Sofia Karina Trommlerová, "Cash Transfers and Fertility: How the Introduction and Cancellation of a Child Benefit Affected Births and Abortions," *Journal of Human Resources*, 2023, *58* (3), 783–818.

Government Accountability Office, *Targeted Jobs Tax Credit: Employer Actions to Recruit, Hire, and Retain Eligible Workers Vary*, US Government Accountability Office, HRD-91-33, 1991.

_ , *Work Opportunity Tax Credit: Employers Do Not Appear to Dismiss Employees to Increase Tax Credits: Report to the Chairman, Subcommittee on Oversight, Committee on Ways and Means, House of Representatives*, US General Accounting Office, 2001.

_ , *Business Tax Incentives: Incentives to Employ Workers With Disabilities Receive Limited Use and Have an Uncertain Impact*, US Government Accountability Office, GAO-03-39, 2002.

Grawe, Nathan, "Lifecycle Bias in Estimates of Intergenerational Earnings Persistence," *Labour economics*, 2006, *13* (5), 551–570.

Groh, Matthew, Nandini Krishnan, David McKenzie, and Tara Vishwanath, "Do Wage Subsidies Provide a Stepping-Stone to Employment for Recent College Graduates? Evidence from a Randomized Experiment in Jordan," *Review of Economics and Statistics*, 2016, *98* (3), 488–502.

Gunderson, Jill and Julie Hotchkiss, "Job Separation Behavior of WOTC Workers: Results from a Unique Case Study," *Social Service Review*, 2007, *81* (2), 317–342.

Hamersma, Sarah, "The Work Opportunity and Welfare-to-Work Tax Credits," Technical Report 2005.

_ , "The Effects of an Employer Subsidy on Employment Outcomes: A Study of the Work Opportunity and Welfare-to-Work Tax Credits," *Journal of Policy Analysis and Management*, 2008, *27* (3), 498–520.

_ , "Why Don't Eligible Firms Claim Hiring Subsidies? The Role of Job Duration," *Economic Inquiry*, 2011, *49* (3), 916–934.

_ **and Carolyn Heinrich**, "Temporary Help Service Firms' Use of Employer Tax Credits: Implications for Disadvantaged Workers' Labor Market Outcomes," *Southern Economic Journal*, 2008, *74* (4), 1123–1148.

Heckman, James and Robert Willis, "Estimation of a Stochastic Model of Reproduction: An Econometric Approach," in "Household production and consumption," National Bureau of Economic Research, 1976, pp. 99–146.

Hnatkovska, Viktoria, Amartya Lahiri, and Sourabh Paul, "Breaking the Caste Barrier Intergenerational Mobility in India," *Journal of Human Resources*, 2013, *48* (2), 435–473.

Hollenbeck, Kevin and Richard Wilke, "The Employment and Earnings Impacts of the Targeted Jobs Tax Credit," Technical Report, W.E. Upjohn Institute for Employment Research, Staff Working Paper 91-07 1991.

Hotz, Joseph and Robert Miller, "Conditional Choice Probabilities and the Estimation of Dynamic Models," *The Review of Economic Studies*, 07 1993, *60* (3), 497–529.

Hunt, Priscilla, Rossana Smart, Lisa Jonsson, and Flavia Tsang, "Breaking Down Barriers: Experiments Into Policies That Might Incentivize Employers to Hire Exoffenders," 2018.

Huttunen, Kristiina, Jukka Pirttilä, and Roope Uusitalo, "The Employment Effects of Low-Wage Subsidies," *Journal of Public Economics*, 2013, *97*, 49–60.

Internal Revenue Service, "Work Opportunity Tax Credit," `https://www.irs.gov/businesses/small-businesses-self-employed/work-opportunity-tax-credit` 2024. Accessed: 2025-01-24.

Iversen, Vegard, Anirudh Krishna, and Kunal Sen, "Rags to Riches? Intergenerational Occupational Mobility in India," *Economic and Political Weekly*, 2017, *52* (44), 107–114.

Jacob, Suraj and Sreeparna Chattopadhyay, "Marriage Dissolution in India: Evidence from Census 2011," *Economic and Political Weekly*, 2016, *51* (33), 25–27.

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer, "Worker Beliefs About Outside Options," *The Quarterly Journal of Economics*, 2024, *139* (3), 1505–1556.

Jayachandran, Seema, "The Roots of Gender Inequality in Developing Countries," *Annual Review of Economics*, 2015, *7* (1), 63–88.

_ , "Fertility decline and Missing Women," *American Economic Journal: Applied Economics*, 2017, *9* (1), 118–139.

_ , "Ten Facts about Son Preference in India," *National Bureau of Economic Research No. 31883*, 2023.

_ **and Rohini Pande**, "Why Are Indian Children So Short? The Role of Birth Order and Son Preference," *American Economic Review*, 2017, *107* (9), 2600–2629.

Jodhka, Surinder, "Sikhism and The Caste Question: Dalits and Their Politics in Contemporary Punjab," *Contributions to Indian Sociology*, 2004, *38* (1-2), 165–192.

Katz, Lawrence, "Wage Subsidies for the Disadvantaged," *National Bureau of Economic Research, No. 5679*, 1996.

Kearney, Melissa and Phillip Levine, "Subsidized Contraception, Fertility, and Sexual Behavior," *The Review of Economics and Statistics*, 2009, *91* (1), 137–151.

Keith-Jennings, Brynne and Raheem Chaudhry, "Most Working-Age SNAP Participants Work, but Often in Unstable Jobs," *Washington, DC: Center on Budget and Policy Priorities*, 2018.

Khanna, Gaurav, "Does Affirmative Action Incentivize Schooling? Evidence from India," *Review of Economics and Statistics*, 2020, *102* (2), 219–233.

Kline, Patrick, Evan Rose, and Christopher Walters, "A Discrimination Report Card," *American Economic Review*, 2024, *114* (8), 2472–2525.

Koenig, Michael, Gillian Foo, and Ketan Joshi, "Quality of Care Within the Indian Family Welfare Programme: A Review of Recent Evidence," *Studies in Family Planning*, 2000, *31* (1), 1–18.

Kundu, Anustup and Kunal Sen, "Multigenerational Mobility in India," *IZA Discussion Paper No. 14566*, 2021.

Laroque, Guy and Bernard Salanié, "Identifying the Response of Fertility to Financial Incentives," *Journal of Applied Econometrics*, 2014, *29* (2), 314–332.

Larsen, Ulla, Woojin Chung, and Monica Dasgupta, "Fertility and Son Preference in Korea," *Population Studies*, 1998, *52* (3), 317–325.

Lee, SangYoon and Ananth Seshadri, "On the Intergenerational Transmission of Economic Status," *Journal of Political Economy*, 2019, *127* (2), 855–921.

Li, Qi and Juan Pantano, "The Demographic Consequences of Sex-Selection Technology," *Quantitative Economics*, 2023, *14* (1), 309–347.

Lorenz, Edward, "TJTC and the Promise and Reality of Redistributive Vouchering and Tax Credit Policy," *Journal of Policy Analysis and Management*, 1995, *14* (2), 270–290.

MacKinnon, James, Morten Nielsen, and Matthew Webb, "Cluster-Robust Inference: A Guide to Empirical Practice," *Journal of Econometrics*, 2023, *232* (2), 272–299.

Mayer, Adrian, "Caste in an Indian Village: Change and Continuity 1954–1992," *Caste Today*, 1996, pp. 32–64.

Mendelsohn, Oliver, "The Transformation of Authority in Rural India," *Modern Asian Studies*, 1993, *27* (4), 805–842.

Merwin, Grier, "Vasectomies: Who's Having Them...and Why?," *Health Education Monographs*, 1974, *2* (3), 260–277.

Milligan, Kevin, "Subsidizing the Stork: New Evidence on Tax Incentives and Fertility," *Review of Economics and Statistics*, 2005, *87* (3), 539–555.

Ministry of Health and Family Welfare, "National Population Policy 2000," Report 2000.

Munshi, Kaivan, "Caste and the Indian Economy," *Journal of Economic Literature*, 2019, *57* (4), 781–834.

Nakao, Keiko and Judith Treas, "Updating Occupational Prestige and Socioeconomic Scores: How the New Measures Measure Up," *Sociological Methodology*, 1994, pp. 1–72.

Neumark, David and Diego Grijalva, "The Employment Effects of State Hiring Credits," *ILR Review*, 2017, *70* (3), 716–749.

Noord-Zaadstra, Boukje, Caspar Looman, Hans Alsbach, Dik Habbema, Egbert te Velde, and Jan Karbaat, "Delaying Childbearing: Effect of Age on Fecundity and Outcome of Pregnancy," *British Medical Journal*, 1991, *302* (6789), 1361–1365.

Norling, Johannes, "Measuring Heterogeneity in Preferences over the Sex of Children," *Journal of Development Economics*, 2018, *135*, 199–221.

Osborne, Evan, "Culture, Development, and Government: Reservations in India," *Economic Development and Cultural Change*, 2001, *49* (3), 659–685.

Oster, Emily, "Hepatitis B and the Case of the Missing Women," *Journal of Political Economy*, 2005, *113* (6), 1163–1216.

Pallais, Amanda, "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review*, 2014, *104* (11), 3565–3599.

Pande, Rohini, "Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India," *American Economic Review*, 2003, *93* (4), 1132–1151.

_ **and Nan Marie Astone**, "Explaining Son Preference in Rural India: The Independent Role of Structural versus Individual Factors," *Population Research and Policy Review*, 2007, *26* (1), 1–29.

Perloff, Jeffrey and Michael Wachter, "The New Jobs Tax Credit: an Evaluation of the 1977-78 Wage Subsidy Program," *American Economic Review*, 1979, *69* (2), 173–179.

Piketty, Thomas, "Theories of Persistent Inequality and Intergenerational Mobility," *Handbook of Income Distribution*, 2000, *1*, 429–476.

Pop-Eleches, Cristian, “The Impact of an Abortion Ban on Socioeconomic Outcomes of Children: Evidence from Romania,” *Journal of Political Economy*, 2006, *114* (4), 744–773.

Prakash, Nishith, “The Impact of Employment Quotas on the Economic Lives of Disadvantaged Minorities in India,” *Journal of Economic Behavior & Organization*, 2020, *180*, 494–509.

Qian, Nancy, “Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earnings on Sex Imbalance,” *The Quarterly Journal of Economics*, 2008, *123* (3), 1251–1285.

_ , “Quantity-Quality and the One Child Policy: The Only-Child Disadvantage in School Enrollment in Rural China,” Working Paper 14973, National Bureau of Economic Research 2009.

Ramaiah, Avatthi, “Identifying Other Backward Classes,” *Economic and Political Weekly*, 1992, pp. 1203–1207.

Ray, Tridip, Arka Roy Chaudhuri, and Komal Sahai, “Whose Education Matters? An Analysis of Inter Caste Marriages in India,” *Journal of Economic Behavior & Organization*, 2020, *176*, 619–633.

Restuccia, Diego and Carlos Urrutia, “Intergenerational Persistence of Earnings: The Role of Early and College Education,” *American Economic Review*, 2004, *94* (5), 1354–1378.

Rosenzweig, Mark and Kenneth Wolpin, “Testing the Quantity-Quality Fertility Model: The Use of Twins as a Natural Experiment,” *Econometrica*, 1980, *48* (1), 227–240.

Roussille, Nina, “The Role of the Ask Gap in Gender Pay Inequality,” *Quarterly Journal of Economics*, 2024, p. qjae004.

Rudner, David, *Caste and Capitalism in Colonial India: The Nattukottai Chettiars.*, Oxford University Press, 1996.

Saez, Emmanuel, Benjamin Schoefer, and David Seim, "Payroll Taxes, Firm Behavior, and Rent Sharing: Evidence from a Young Workers' Tax Cut in Sweden," *American Economic Review*, 2019, *109* (5), 1717–1763.

Salve, Pradeep and Chander Shekhar, "Disappearing Male Sterilization in India: Do We Care?," *Contraception and Reproductive Medicine*, 2023, *8* (1), 31.

Schünemann, Benjamin, Michael Lechner, and Conny Wunsch, "Do Long-Term Unemployed Workers Benefit from Targeted Wage Subsidies?," *German Economic Review*, 2015, *16* (1), 43–64.

Sekher, T.V., "Special Financial Incentive Schemes for the Girl Child in India: A Review of Select Schemes," `https://www.unfpa.org/sites/default/files/resource-pdf/UNFPA_Publication-39772.pdf` 2010.

_ , "Ladlis and Lakshmis: Financial Incentive Schemes for the Girl Child," *Economic and Political Weekly*, 2012, pp. 58–65.

_ **and F Ram**, "Conditional Cash Transfers for Girls in India: Assessment of a Girl Child Promotion Scheme from Beneficiary Perspective," Technical Report, International Institute for Population Sciences 2015.

Sen, Amartya, "Missing Women," *British Medical Journal*, 1992, *304* (6827), 587.

_ , "Missing Women—Revisited," *British Medical Journal*, 2003, *327* (7427), 1297–1298.

Sinha, Nistha and Joanne Yoong, "Long-Term Financial Incentives and Investment in Daughters: Evidence From Conditional Cash Transfers In North India," *World Bank Working Paper No. 4860*, 2009.

Solon, Gary, "Intergenerational Mobility in the Labor Market," *Handbook of Labor Economics*, 1999, *3*, 1761–1800.

Todd, Petra and Kenneth Wolpin, "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 2006, *96* (5), 1384–1417.

U.S. Department of Labor Office of Inspector General, "Targeted Jobs Tax Credit: Employment Inducement or Employer Windfall?," https://www.oig.dol.gov/public/reports/oa/pre_1998/04-94-021-03-320s.htm 1994. Accessed: 2025-01-24.

Walker, James, "The Effect of Public Policies on Recent Swedish Fertility Behavior," *Journal of Population Economics*, 1995, *8* (3), 223–251.

Willis, Robert, "A New Approach to the Economic Theory of Fertility Behavior," *Journal of Political Economy*, 1973, *81* (2, Part 2), S14–S64.

Wooldridge, Jeffrey, *Econometric Analysis of Cross Section and Panel Data*, 2 ed., MIT Press, 2010.

Zimmerman, David, "Regression Toward Mediocrity in Economic Stature," *American Economic Review*, 1992, pp. 409–429.

Appendix A: Chapter 1

A.1 Additional Figures

Figure A1: Number of Children for Women with Completed Fertility

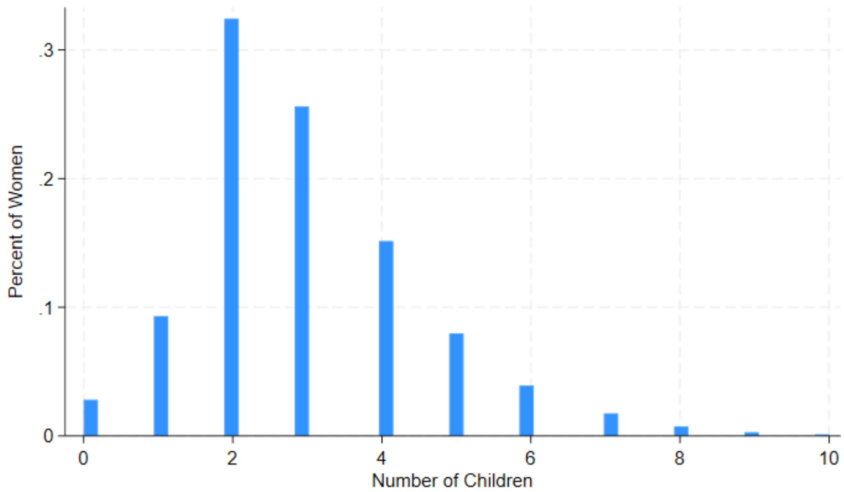

Note: This figure shows the histogram of the number of children for women aged 30+ in 2015-16. Data source: DHS

Figure A2: Effect of Girl-Boy Policy on Having Two Children For Women with Two Children Pre-Policy

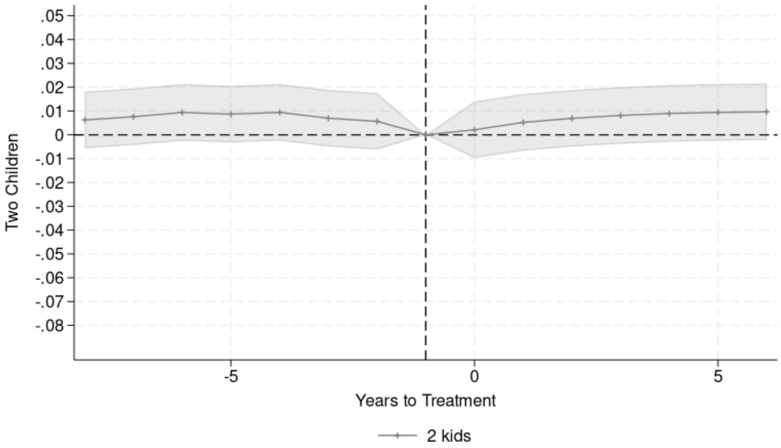

Note: This figure shows the effect of the Girl-Boy policy on the probability that a woman has two kids, conditional on having two kids in the first year of policy implementation, i.e., not continuing childbearing after the policy. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure A3: Effect of Polices on Subsidy Eligible Compositions: Synthetic Control Method

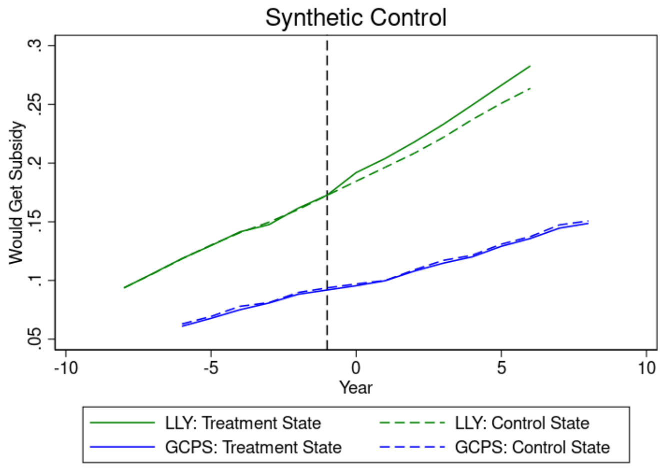

Note: This figure presents synthetic control estimates of the effect of the Girl-Boy and Girls-Only policies on the probability that a woman has a subsidy-eligible child composition. Each panel compares the treated state to a synthetic control constructed from untreated states matched on pre-policy trends.

Figure A4: Effect of Policies on Subsidy-Eligible Child Compositions by Age at Time of Policy Implementation: Synthetic Control Method

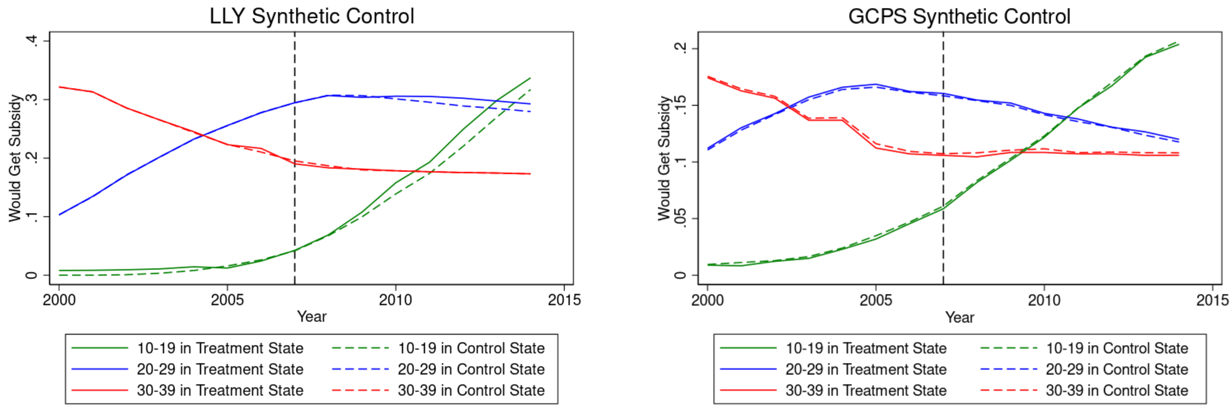

(a) Girl-Boy Policy (b) Girls-Only Policy

Note: This figure presents synthetic control estimates of the effect of the Girl-Boy and Girls-Only policies on the probability that a woman has a subsidy-eligible child composition, stratified by age at the time of policy implementation. Each panel compares the treated state to a synthetic control constructed from untreated states matched on pre-policy trends.

Figure A5: Effect of Policies on Sterilization at Non-Subsidy-Eligible Compositions

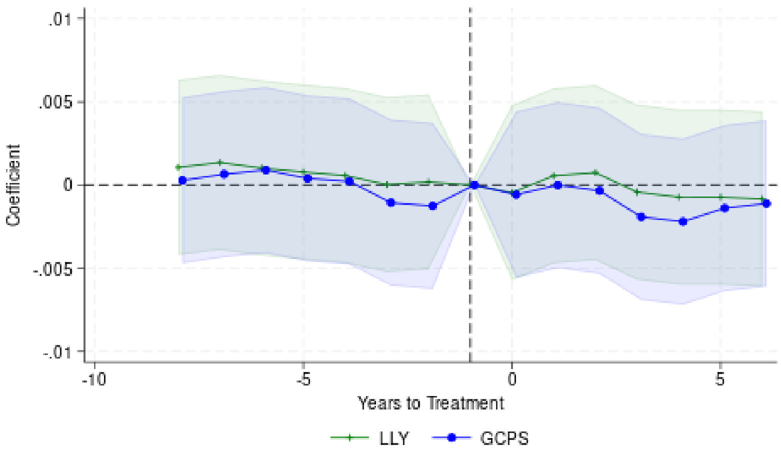

Note: This figure shows the effect of the Girl-Boy (green) and Girls-Only (green) policies on the probability of sterilization at a subsidy-ineligible child composition. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure A6: Effect of Policies on Sterilization at Subsidy-Eligible Compositions Within Policy-Prescribed Time

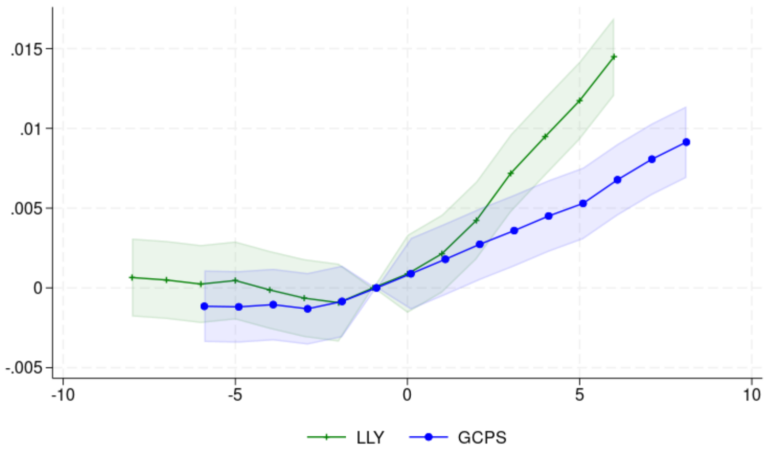

Note: This figure shows the effect of the Girls-Boy (green) and Girls-Only (blue) policy on the probability of sterilization at a subsidy-eligible child composition within the policy prescribed timeframe. The figure plots the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Each regression includes fixed effects for states, years, and age of the woman in the first year of the policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure A7: Effect of Policies on Sterilization at Subsidy-Eligible Compositions Within Policy-Prescribed Time

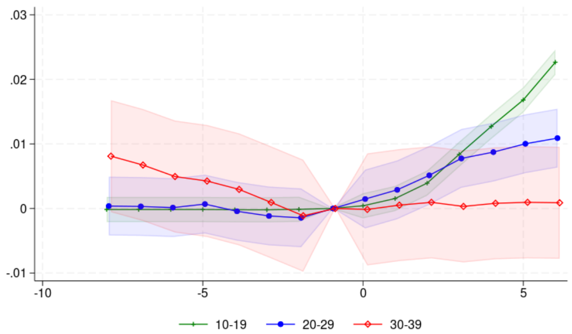

(a) Girl-Boy Policy

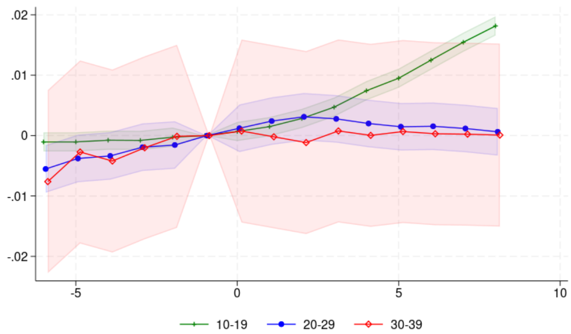

(b) Girls-Only Policy

Note: This figure shows the effect of the Girl-Boy and Girls-Only policies on the probability of sterilization at a subsidy-eligible child composition within the policy-prescribed time frame, stratified by age group of the woman in the first year of policy implementation. Each plot displays the β_z coefficient and its 95% confidence interval derived from estimating the event study equation (1.2). Regressions include state, year, and woman's age fixed effects, with covariates for religion, caste, and urban/rural status. Standard errors are clustered at the state level using wild cluster bootstrap.

Figure A8: Effect of Policies on Mortality Outcomes of Policy-Eligible Girls

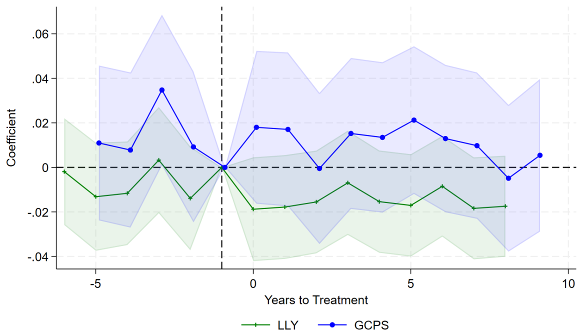

(a) Infant Mortality (Under Age 1)

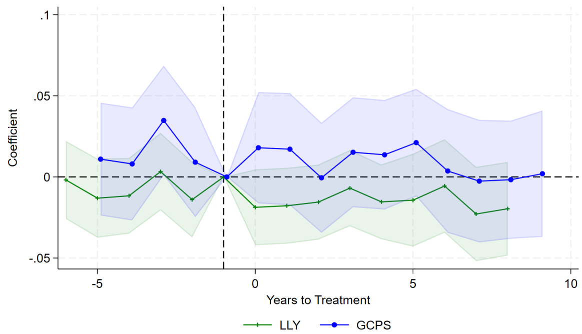

(b) Under-5 Mortality

Note: This figure shows the effect of the Girl-Boy (green) and Girls-Only (blue) policies on mortality outcomes for policy-eligible girls. The plots display the β_z coefficients and 95% confidence intervals estimated from the event study equation (1.2). Each regression includes state and birth year fixed effects, along with controls for religion, caste, and urban/rural status. Standard errors are clustered at the state level using wild cluster bootstrap.

A.2 Additional Tables

Appendix Tables

Table A1: Districts with Dhanlaxmi Policy

State	**District**	**Block**
Andhra Pradesh	Khammam	Aswaraopeta
	Warangal	Narsampet
Chattisgarh	Bastar	Jagdalpur
	Bijapur	Bhopalpattnam
Orissa	Malkangiri	Kalimela
	Koraput	Semiliguda
Jharkhand	Giridih	Tisri
	Kodarma	Markachor
Bihar	Jamoi	Sono
Uttar Pradesh	Rae Bareilly	Shivgarh
Punjab	Fatehgarh Sahib	Sirhind

*Note:*This table shows the list of districts that are excluded from the analysis owing to the implementation of a similar girl-promotion CCT called Dhanlaxmi.

Table A2: Effect of Polices on Subsidy Eligible Compositions: Bordering Districts Analysis

	Girl-Boy Policy				Girls-Only Policy			
	All	10–19	20–29	30–39	All	10–19	20–29	30–39
Treat x Post	0.0199***	0.0264***	0.0185***	0.0069	0.0032	0.0031	0.0033	−0.0004
	(0.0021)	(0.0024)	(0.0037)	(0.0051)	(0.0045)	(0.0058)	(0.0076)	(0.0100)
Mean	0.1733	0.0841	0.2476	0.1873	0.1065	0.0844	0.1311	0.0943
Observations	959,220	365,340	414,630	179,250	173,640	64,320	75,030	34,290

*Note:*This table shows the point estimate and the standard error of the effect of the the Girl-Boy and Girls-Only policies on the probability of having a subsidy-eligible child composition using the border strategy. Each coefficient is estimated using the difference-in-differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the district level.

Table A3: Effect of Policies on Sterilization at Subsidy-Eligible Compositions Within Policy-Prescribed Time (DiD estimates)

	Girl-Boy	Girls-Only
Treat × Post	0.0071***	0.0063***
	(0.0005)	(0.0004)
Mean (Control States)	0.0206	0.0077
Observations	5,130,855	2,234,685

Note: This table shows the point estimate and the standard error of the effect of the Girl-Boy and girls-only policies on the probability of a woman sterilizing at a subsidy-eligible child composition within the policy prescribed time. Each regression is estimated using the difference in differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

Table A4: Effect of Policies on Timely Sterilization at Subsidy-Eligible Child Compositions By Age of Woman at Time of Policy Implementation (DiD Estimates)

	Girl-Boy Policy			Girls-Only Policy		
	10–19	20–29	30–39	10–19	20–29	30–39
Treat × Post	0.0096***	0.0069***	−0.0027	0.0087***	0.0045***	0.0029
	(0.0004)	(0.0008)	(0.0017)	(0.0003)	(0.0007)	(0.0029)
Mean (Control States)	0.0053	0.0397	0.0532	0.0015	0.0111	0.0203
Observations	2,200,425	2,208,960	721,470	1,017,255	1,001,460	215,970

Note: This table presents the effect of the Girl-Boy and Girls-Only policies on the probability that a woman sterilizes at a subsidy-eligible child composition within the policy-prescribed time-frame, stratified by age group at policy implementation. Estimates are based on the difference-in-differences specification in equation (1.1), including fixed effects for state, year, and age, and controlling for religion, caste, and urban/rural status. Standard errors are clustered at the state level using wild cluster bootstrap.

Table A5: Effect of Policies on Infant and Under-5 Mortality Rates (DiD Estimates)

	Infant Mortality Rate (IMR)		Under-5 Mortality Rate (U5MR)	
	Girl-Boy	Girls-Only	Girl-Boy	Girls-Only
Treat × Post	−0.0070	−0.0018	−0.0072	−0.0027
	(0.0043)	(0.0068)	(0.0047)	(0.0070)
Mean (Control States)	0.0795	0.0469	0.0798	0.0471
Observations	111,863	48,705	97,162	43,321

*Note:*This table shows the point estimate and the standard error of the effect of the the Girl-Boy and Girls-Only policies on the infant and under-5 mortality rates of policy-eligible girls. Each coefficient is estimated using the difference-in-differences equation (1.1) and includes fixed effects for states, years, and age of woman in the first year of policy implementation. Covariates such as religion, caste, urban/rural status are also included. Standard errors are clustered at the state level using wild cluster bootstrap.

A.3 Methods Appendix

A.3.1 Control States in Difference-in-Differences Strategy

Girl-Boy Policy I begin by constructing a set of nearby control states for Madhya Pradesh (the treated state under the Girl-Boy policy). I define "nearby" as all Indian states (excluding Union Territories due to small sample sizes) within a 750 km radius of the centroid of Madhya Pradesh (denoted by the green circle in Figure A9), as well as any state outside this radius that is not contiguous but has at most one intervening state. This yields Rajasthan, Gujarat, Maharashtra, Karnataka, Andhra Pradesh, Chhattisgarh, Jharkhand, Odisha, Bihar, Uttar Pradesh, Haryana, and Uttarakhand, plus two states outside the 750 km radius: Punjab and Himachal Pradesh.

Several of these states introduced girl-promotion conditional cash transfer programs during the study period: Karnataka (Bhagyalakshmi, 2006), Andhra Pradesh (Girl Child Protection Scheme, 2005), Jharkhand (Mukhyamantri Ladli Laxmi Yojana, 2011), Haryana (Ladli, 2005), Uttarakhand (Nanda Devi Kanya Dhan Yojana, 2009), Himachal Pradesh (Indira Gandhi Balika Suraksha Yojana, 2007), and Punjab (Balri Rakshak Yojana, 2005). I exclude these states, leaving Rajasthan, Gujarat, Maharashtra, Odisha, Bihar, Uttar Pradesh, and Chhattisgarh as candidate controls.

With this set, I estimate the event-study specification in equation (1.2) using two outcomes: the subsidy-eligible composition of births and sterilization conditional on subsidy-eligible composition. For these outcomes, the joint null that all pre-treatment coefficients equal zero cannot be rejected at the 5% level once Maharashtra is excluded. I therefore retain Bihar, Chhattisgarh, Gujarat, Odisha, Rajasthan, and Uttar Pradesh as the control states for the Girl-Boy policy.

Girls-Only Policy I next construct a set of nearby control states for Andhra Pradesh (the treated state under the Girl-Only policy). Using the same definition of "nearby," I identify Gujarat, Maharashtra, Karnataka, Madhya Pradesh, Chhattisgarh, Odisha, Tamil Nadu, Kerala, and Goa, as well as three states outside the 750 km radius (denoted by the blue circle in Figure A9): Uttar Pradesh, Jharkhand and West Bengal.

Several of these states implemented girl-promotion conditional cash transfer programs during the study period: Karnataka (Bhagyalakshmi, 2006), Madhya Pradesh (Ladli, 2007), Jharkhand (Mukhyamantri Ladli Laxmi Yojana, 2011), Goa (Ladli, 2012), West Bengal (Kanyashree Prakalpa, 2013), and Tamil Nadu (Girl Child Protection Scheme, 1992). I exclude these states, leaving Uttar Pradesh, Chhattisgarh, Gujarat, Kerala, Maharashtra, and Odisha as candidate controls.

With this set, I again estimate the event-study specification in equation (1.2) using the same two outcomes. The joint null that all pre-treatment coefficients equal zero cannot be rejected at the 5% level once Uttar Pradesh is excluded. I therefore use Chhattisgarh, Gujarat, Kerala, Maharashtra, and Odisha as the control states for the Girl-Only policy.

Figure A9: Proximity of States to Treated States

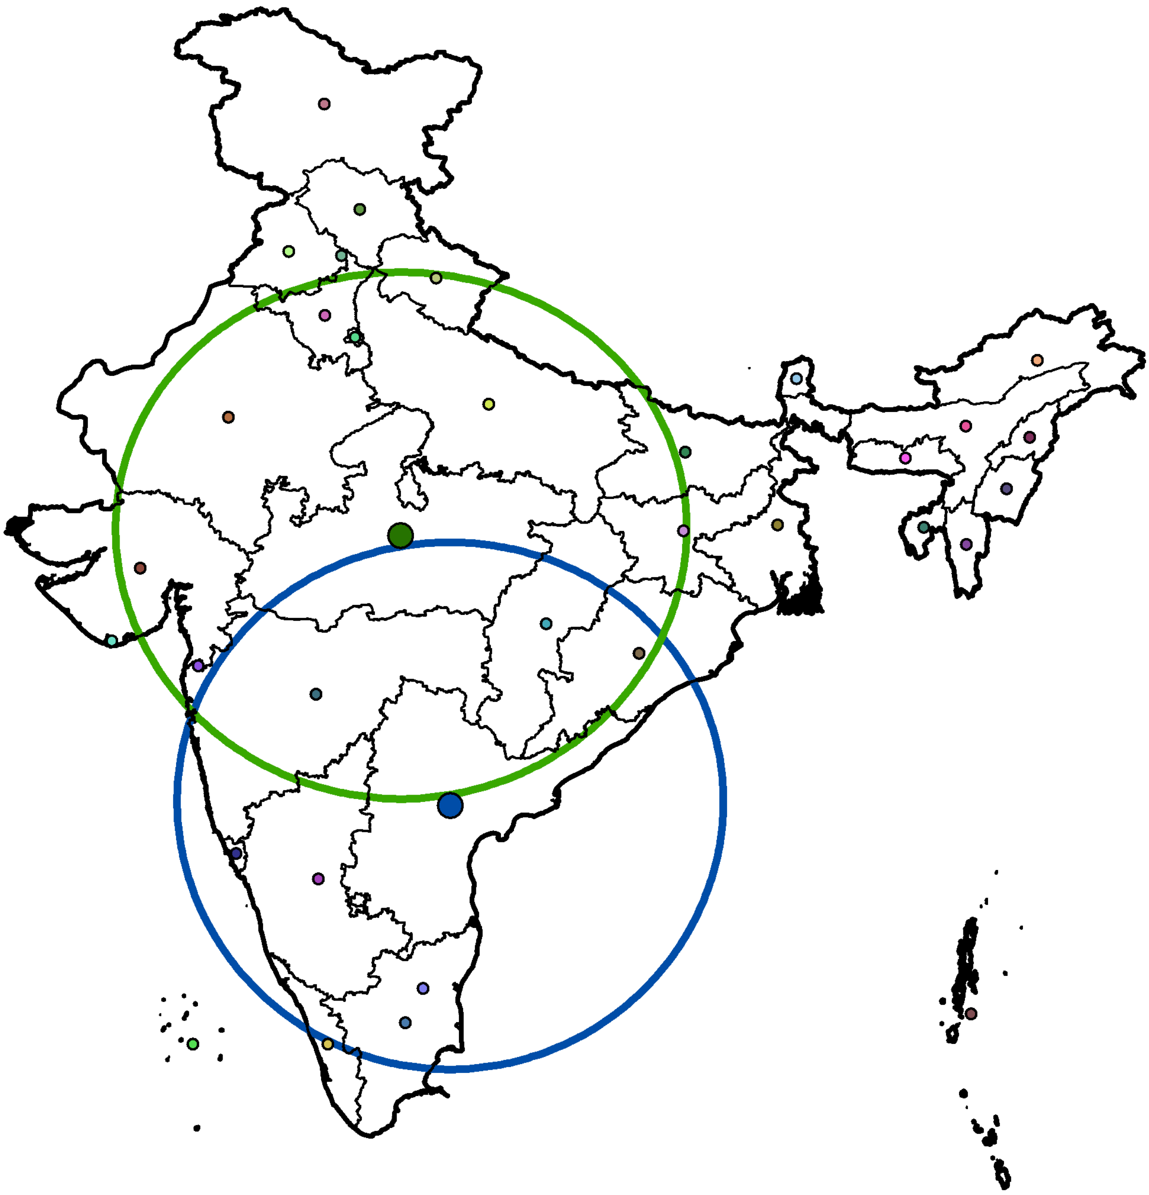

Notes: The figure shows states within a 750 km radius of the centroid of the treated state in the Girl-Boy policy (green) and the Girls-Only policy (blue).

A.3.2 Border Strategy

The state of Madhya Pradesh (the treated state in the Girl-Boy policy) shares its boundary with five Indian states: Gujarat, Rajasthan, Chhattisgarh, Maharashtra, and Uttar Pradesh. 35 districts in Madhya Pradesh bordered at least one of these states: Agar Malwa, Alirajpur, Anuppur, Ashoknagar, Balaghat, Barwani, Betul, Bhind, Burhanpur, Chhatarpur, Chhindwara, Datia, Dindori, Guna, Gwalior, Jhabua, Khandwa, Mandla, Mandsaur, Morena, Neemuch, Panna, Rajgarh, Ratlam, Rewa, Sagar, Satna, Seoni, Shahdol, Sheopur, Shivpuri, Sidhi, Singrauli, Tikamgarh, and Umaria. These districts constitute the treatment group in the border analysis for evaluating the girl-boy policy. The control group consists of neighboring districts in adjoining states. In Maharashtra, the bordering districts are Amravati, Bhandara, Buldhana, Dhule, Gondia, Jalgaon, Nagpur, and Nandurbar. In Rajasthan, the controls are Baran, Banswara, Bhilwara, Chittorgarh, Dholpur, Jhalawar, Karauli, Kota, Pratapgarh, and Sawai Madhopur. In Chhattisgarh, the bordering districts are Balrampur, Bilaspur, Kabirdham, Koriya, Mungeli, Rajnandgaon, and Surajpur. In Gujarat, the controls are Chhota Udaipur and Dahod. Finally, in Uttar Pradesh, the bordering districts are Agra, Allahabad, Auraiya, Banda, Chitrakoot, Etawah, Jalaun, Jhansi, Mahoba, Mirzapur, and Sonbhadra.

The state of Andhra Pradesh, including the newly formed Telangana, served as the treated state for the Girls-Only policy. Andhra Pradesh shares borders with Maharashtra, Tamil Nadu, Odisha, Chhattisgarh, and Karnataka. I exclude Karnataka because it implemented the Bhagyalaxmi program, which provided subsidies for up to two girls in families with three or fewer children, and Tamil Nadu because it operated a girl-promotion conditional cash transfer program of the same design as Andhra Pradesh's scheme (the Girl Child Protection Scheme) during at least some years between 2000 and 2015. The following districts in Andhra Pradesh share a border with Maharashtra, Odisha, and/or Chhattisgarh: Adilabad, East Godavari, Karimnagar, Khammam, Nizamabad, Srikakulam, Visakhapatnam, Vizianagaram, and Warangal. These districts constitute the treatment group for the Girls-Only policy. The control group consists of neighboring districts in Maharashtra, Odisha, and Chhattisgarh. In Maharashtra, the bordering districts are Chandrapur, Gadchiroli, Nanded, and Yavatmal. In Odisha, the bordering districts are Gajapati, Ganjam, Koraput, Malkangiri, and Rayagada. In Chhattisgarh, the bordering districts are Bijapur and Sukma.

Appendix B: Chapter 2

B.1 Additional Figures

Figure B1: Example WOTC subsidy schedule

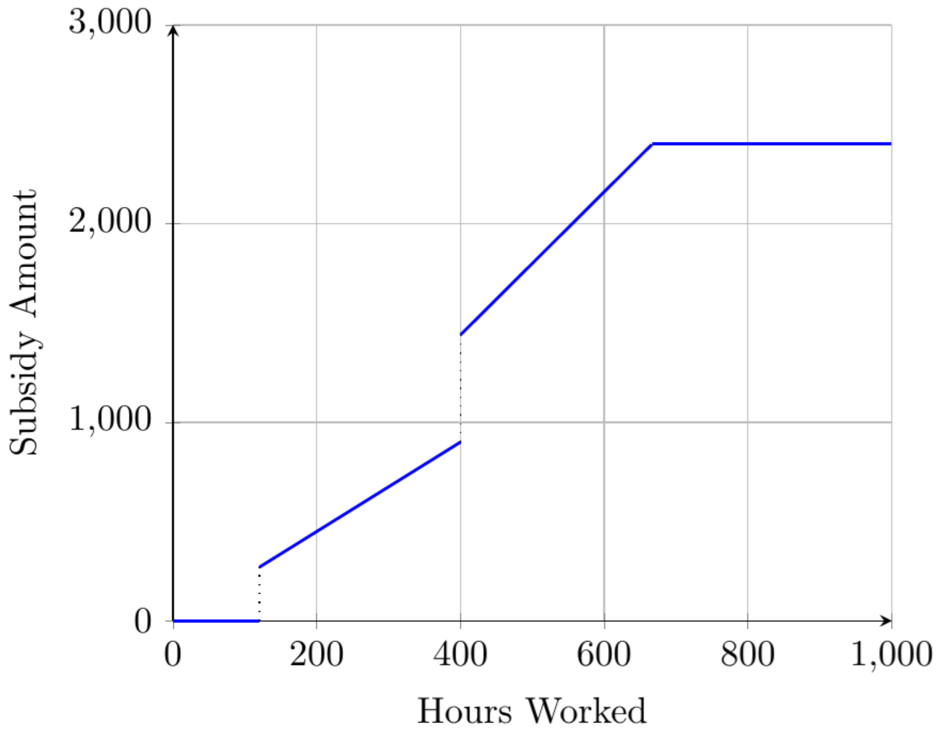

Note: Figure reports the subsidy amount by hours worked for the median worker earning $9 per hour. For the first 120 hours, the subsidy does not pay anything. The subsidy then pays 25% of wages for hours 120 to 400 and 40% of wages for more than 400 hours, up to a cap of $2400.

Figure B2: Hourly Starting Wage

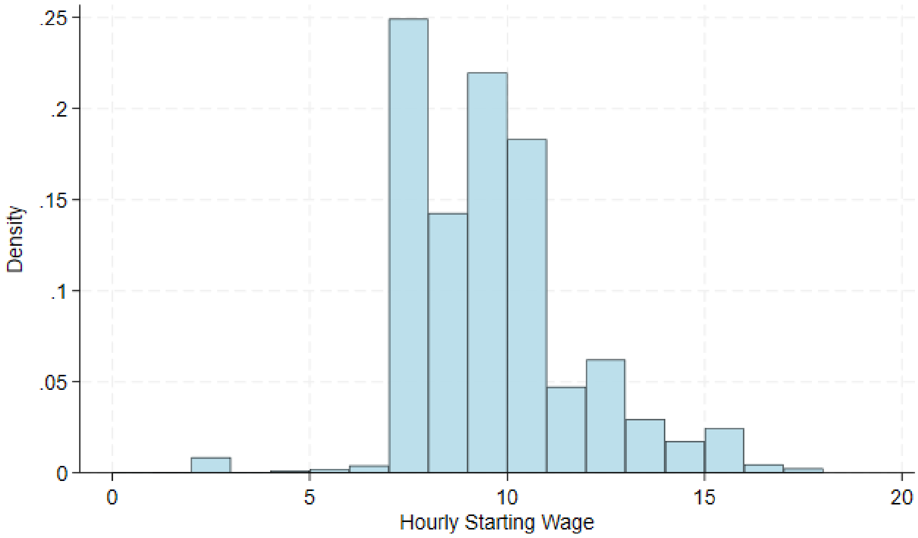

Note: Figure reports the distribution of starting hourly wage for workers in WOTC-certified jobs in Wisconsin between 2009-2020.

Figure B3: Changes in WOTC over time

(a) Hires eligible for WOTC in Wisconsin

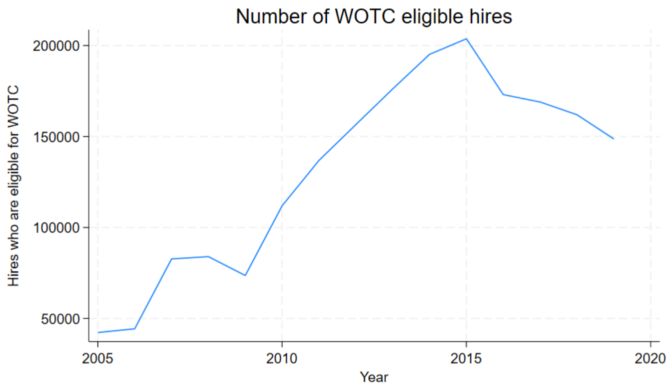

(b) Firm application behavior in Wisconsin

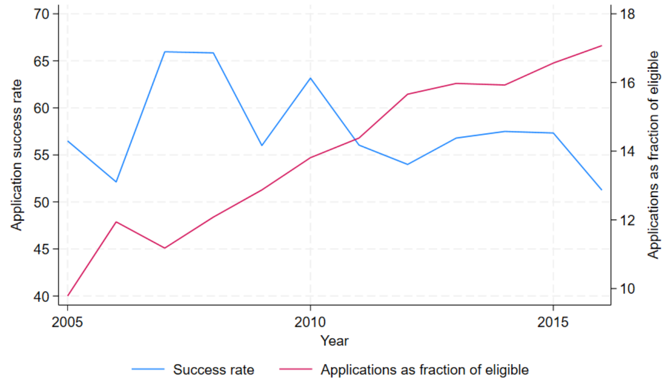

(c) Share of below poverty line individuals eligible for WOTC

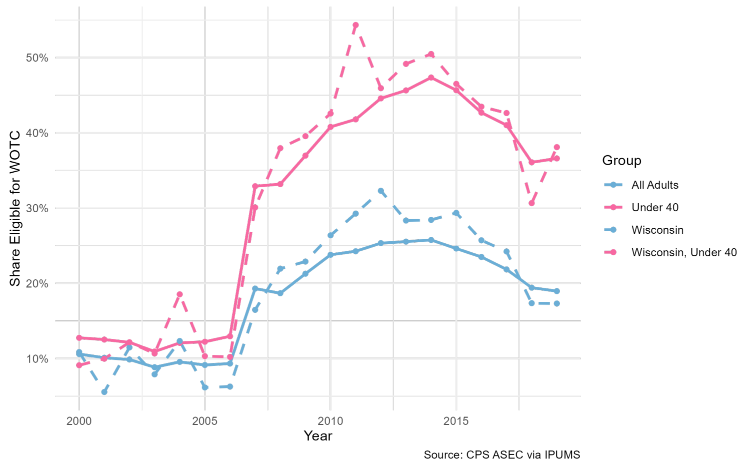

Note: Panel (a) reports the total number of new hires in Wisconsin in a given year for whom our administrative data indicates they are eligible for WOTC due to participation in SNAP or TANF, or conviction for a felony. Panel (b) reports the fraction of WOTC applications that are successful as well as the fraction of eligible individuals for whom firms apply for WOTC. Panel (c) uses data from the Current Population Survey to estimate the fraction of individuals below the poverty line who would be eligible for WOTC if hired into a new job. It estimates this separately for all adults, adults under age 40, and Wisconsin. Estimates are based on CPS data on unemployment, veteran status, and receipt of SNAP, TANF, and SSI. This leaves out those who may be eligible for other reasons (e.g. felony conviction in the previous year). Combined, these omitted categories typically account for around 10% of certifications, meaning the estimates in the figure are likely lower bounds on the true eligibility fraction.

Figure B4: Average Quarterly WOTC applications by State (Per 1000 residents)

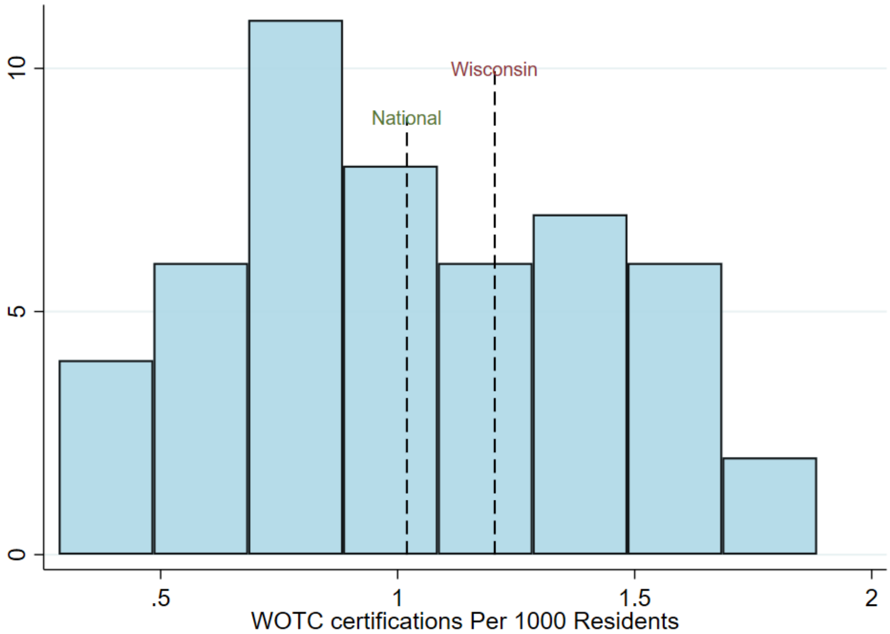

Note: Figure reports average quarterly WOTC applications per 1000 residents by state. The vertical dashed lines indicates the average applications in Wisconsin as well as averaged across all 50 states. Data is taken from the Department of Labor national Work Opportunity Tax Credit data, which is based on quarterly data submitted by state workforce agencies through the the web-based Tax Credit Reporting System.

Figure B5: Effect of WOTC on individual outcomes, event study analysis of SNAP expansion (age 25 sample)

(a) New hire

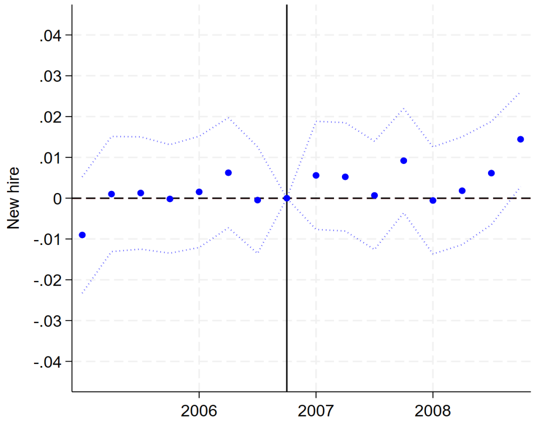

(b) Employed

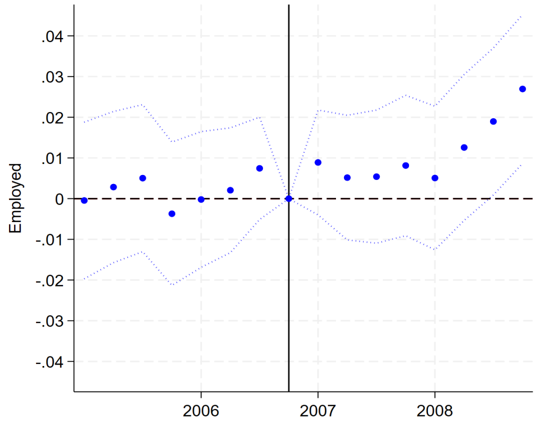

(c) Earnings

(d) New hire into WOTC certified job

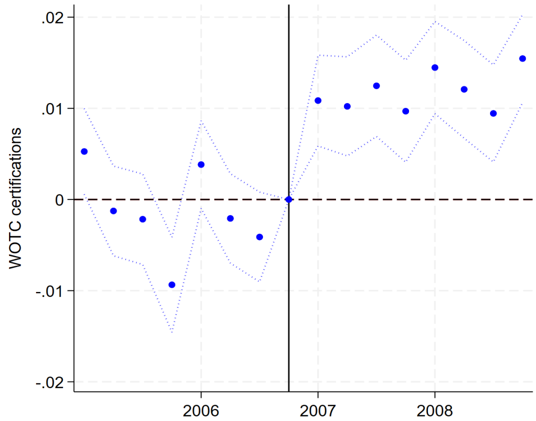

(e) Public benefits

(f) Criminal conviction

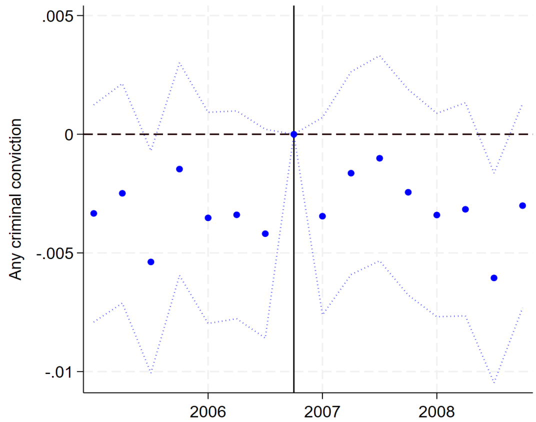

Note: Figure reports event study estimates from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.1) for new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Sample is individuals who are around the age 25 cutoff. Controls include monthly age fixed effects and quarter-year fixed effects. Standard errors clustered at the individual level.

Figure B6: Effect of WOTC on individual outcomes, event study analysis of SNAP expansion (age 40 sample)

(a) New hire

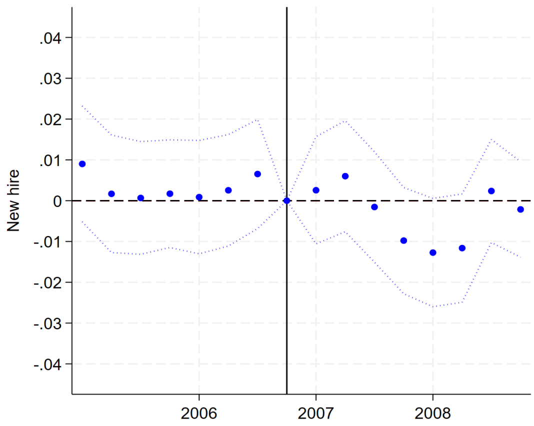

(b) Employed

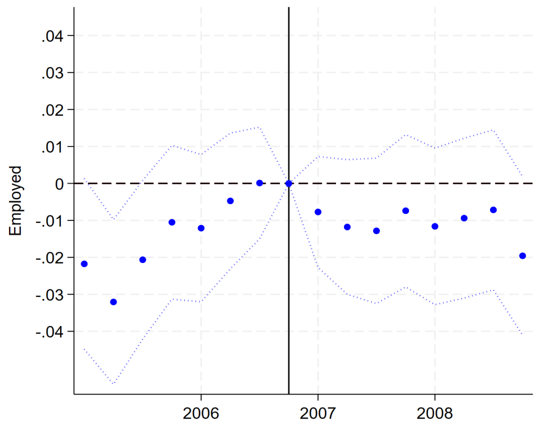

(c) Earnings

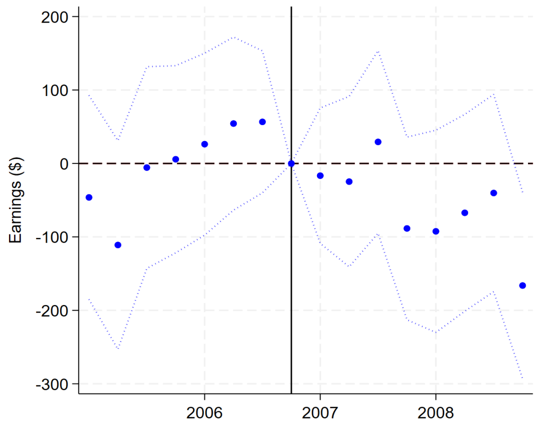

(d) New hire into WOTC certified job

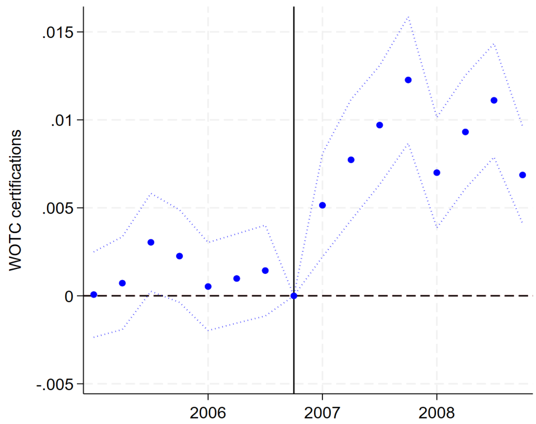

(e) Public benefits

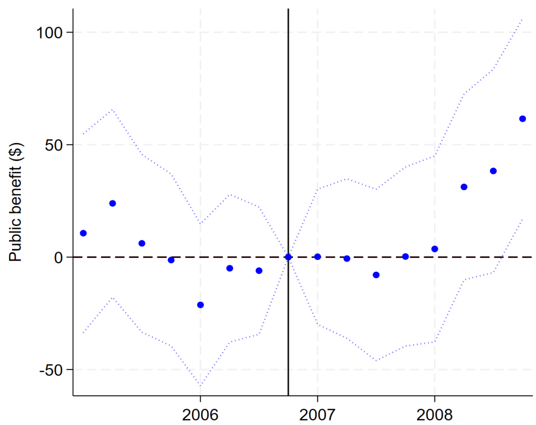

(f) Criminal conviction

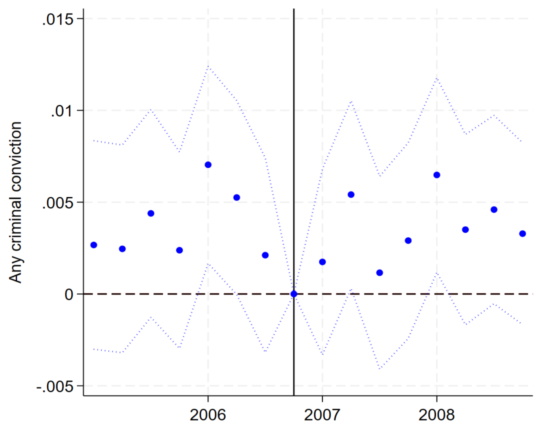

Note: Figure reports event study estimates from the 2007 expansion to SNAP recipients aged 40-39 in Equation (2.1) for new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Sample is individuals who are around the age 40 cutoff. Controls include monthly age fixed effects and quarter-year fixed effects. Standard errors clustered at the individual level.

Figure B7: Effect of WOTC on individual outcomes, event study analysis of SNAP expansion (unemployed sample)

(a) New hire

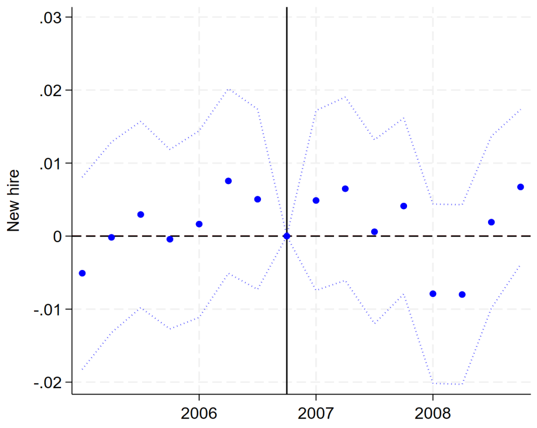

(b) Employed

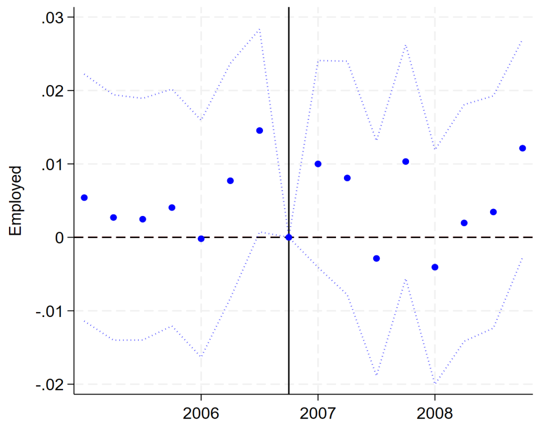

(c) Earnings

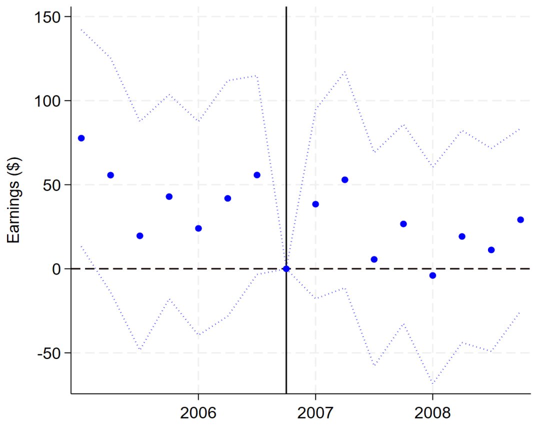

(d) New hire into WOTC certified job

(e) Public benefits

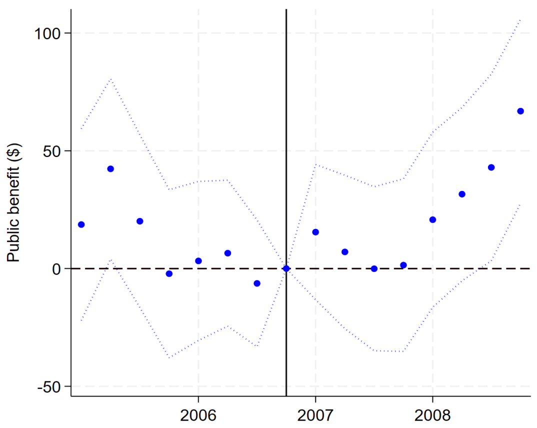

(f) Criminal conviction

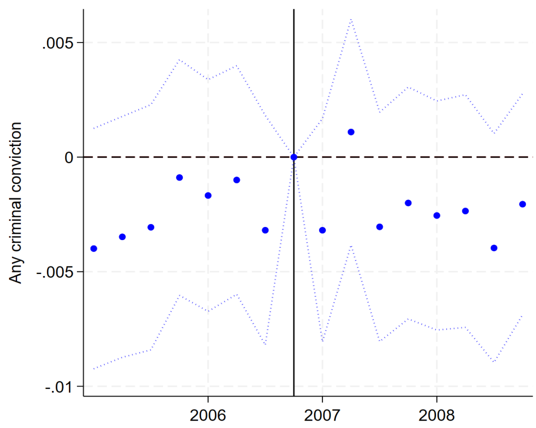

Note: Figure reports stacked event study estimates from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.1) for new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Sample is individuals who entered quarter t unemployed. Controls include monthly age fixed effects and quarter-year by age sample fixed effects. Standard errors clustered at the individual level.

Figure B8: Distribution of the running variable, RD analysis

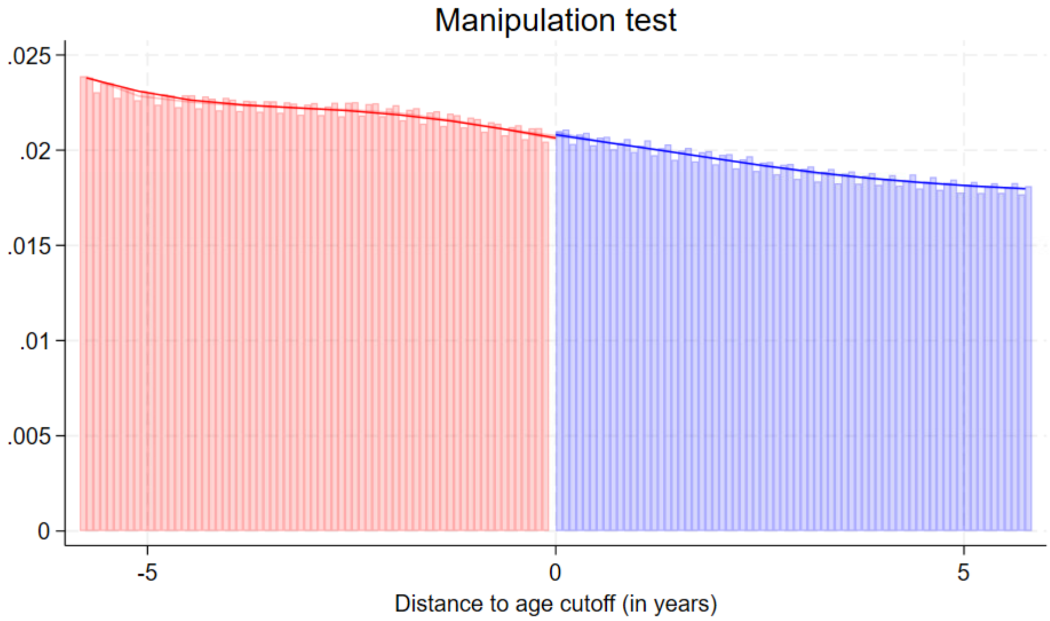

Note: This figure plots the distribution of age around the relevant thresholds and a non-parametric regression for each half of the distribution in order to test for a discontinuity in population density around the threshold (Cattaneo et al., 2020).

Figure B9: Balance on pre-determined characteristics, RD analysis

(a) Employed (t-1)

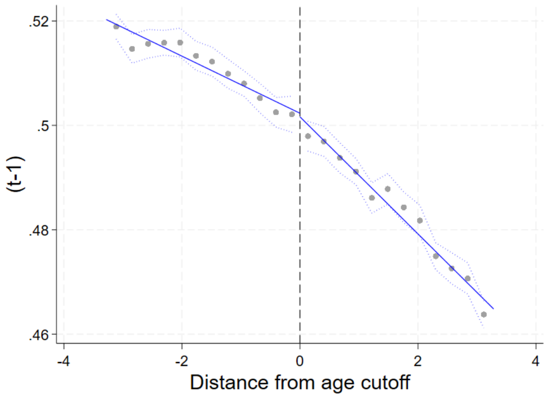

(b) Employed (t-2)

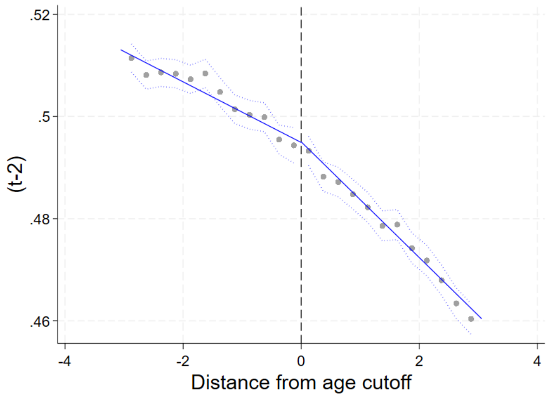

(c) Earnings (t-1)

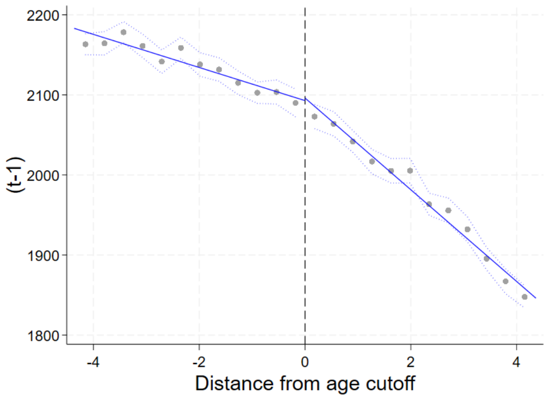

(d) Earnings (t-2)

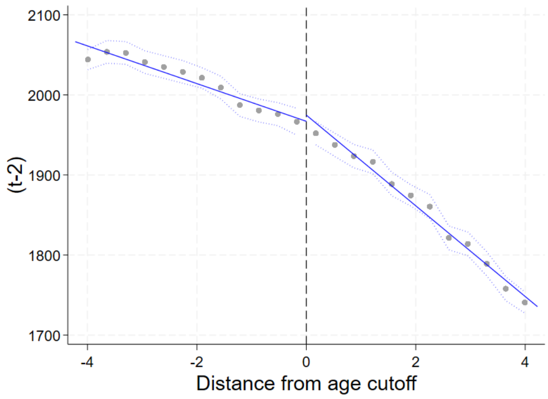

(e) Female

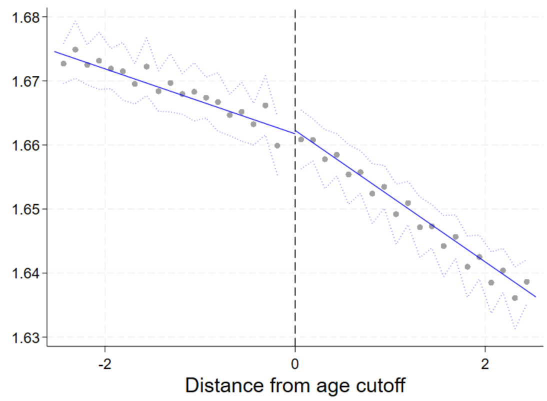

(f) Black

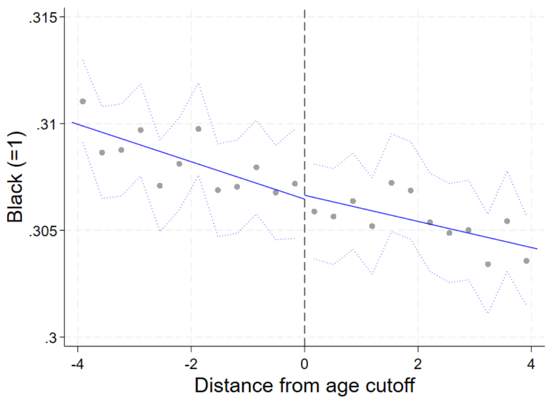

Note: These figures plot the relationship between various characteristics of the individual in the period prior to t and the running variable in order to test for pre-existing imbalances.

Figure B10: Effect of WOTC on individual outcomes, RD analysis (unemployed sample)

(a) Employed

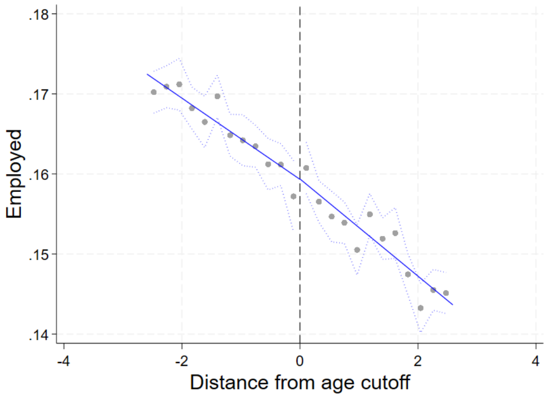

(b) New hire

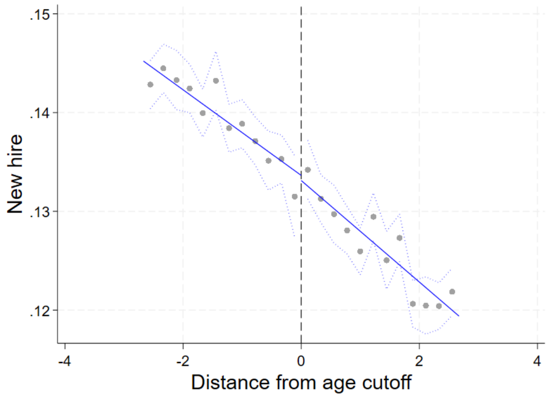

(c) Earnings

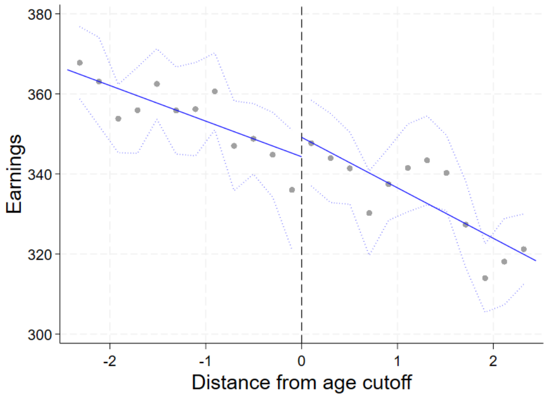

(d) WOTC certified job

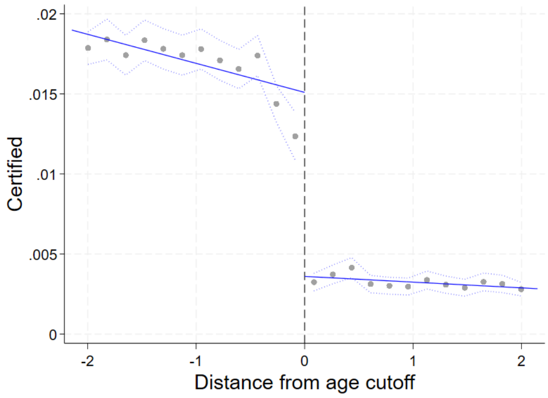

(e) Public benefits

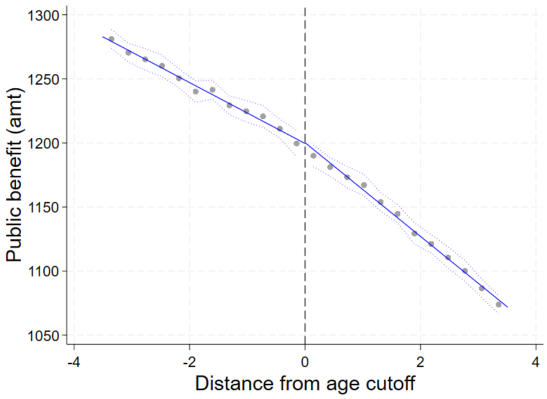

(f) Criminal conviction

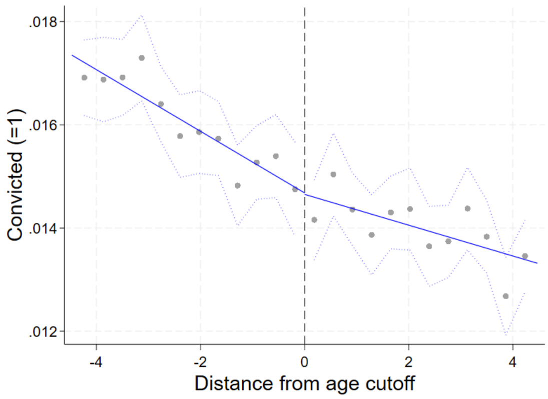

Note: These figures plot the relationship between individual-level outcomes and the difference between the individual's age and the relevant WOTC age eligibility threshold. The data is collected at the quarterly level, and the sample is restricted to individuals who are not employed at the start of the quarter. Individuals whose age is to the left of the threshold are eligible for WOTC, while those to the right of the threshold are not. The bandwidth is based on the optimal bandwidth selection from Cattaneo et al. (2020). Panel (a) examines whether the individual is employed, panel (b) focuses on whether the individual is hired into a new job in that quarter, and panel (c) measures individual earnings in the quarter. Panel (d) plots the first stage relationship between the individual's age and whether they are working in a job with an associated WOTC certification. Panel (e) measures the relationship with the dollar value of three forms of public benefits in a quarter (TANF, Unemployment Insurance, and SNAP), while panel (f) measures if the individual is convicted of a criminal offense in the relevant quarter.

Figure B11: Bootstrap estimates of the ratio of effects on employment and WOTC take-up

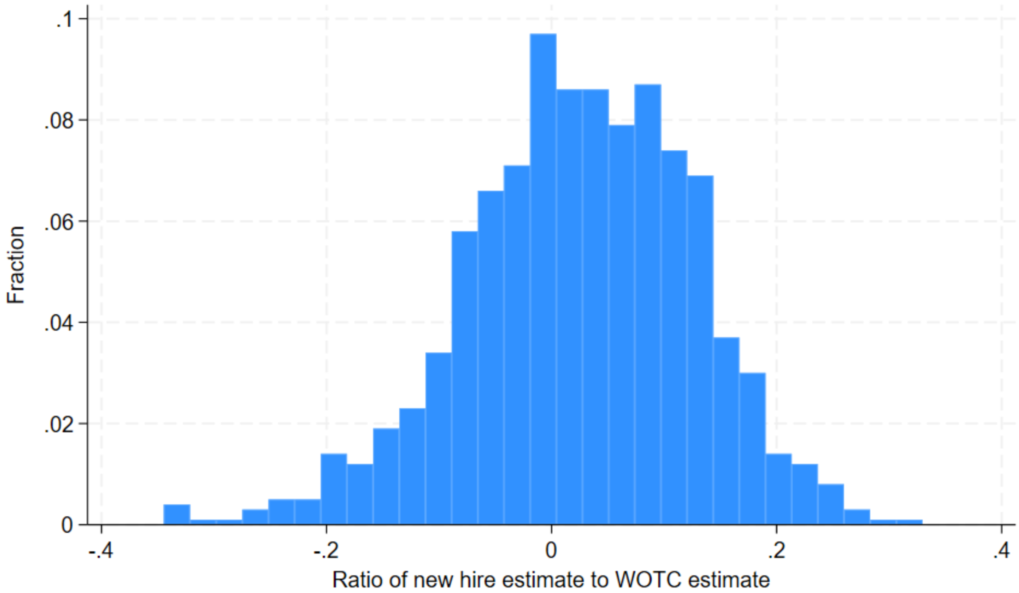

Note: This figure plots the distribution of estimates of the estimated ratio of effects on employment and effects on WOTC take-up. We construct the estimates by drawing block bootstrap samples at the level of the individual, and estimating effects on employment and WOTC uptake within the three individual-level estimation approaches (SNAP expansion DID, Age-based RDD, eWOTC expansion DID). Within each sample and estimation strategy, we estimate the ratio of effects and construct a precision-weighted average of the three estimates. This figure plots the precision-weighted estimates across the bootstrap samples.

Figure B12: Histograms of total earnings for WOTC-certified workers

(a) Top 50 WOTC firm

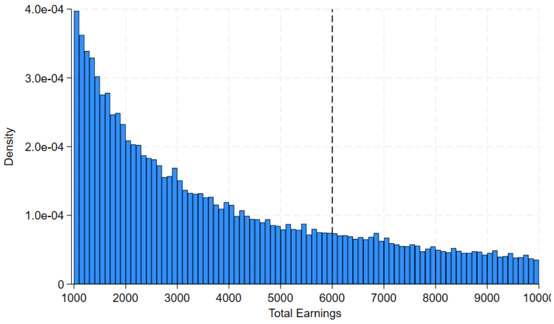

(b) WOTC high success firm

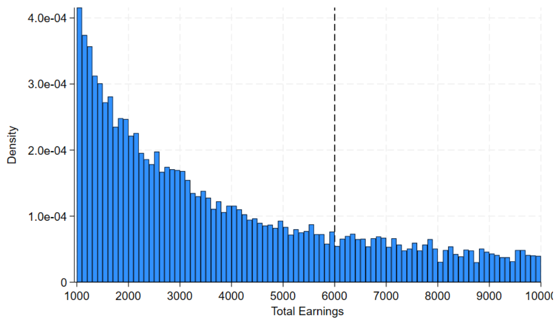

(c) Staffing firm

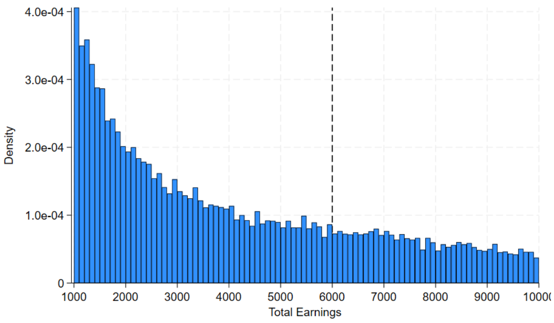

(d) 2007 to 2013

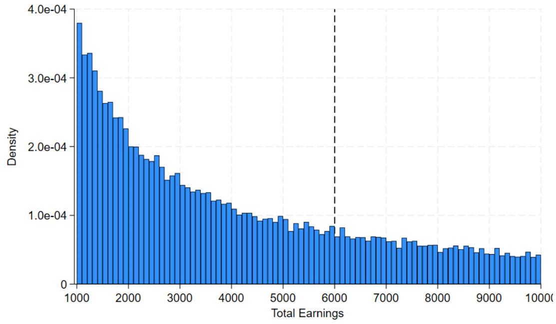

(e) 2014 to 2019

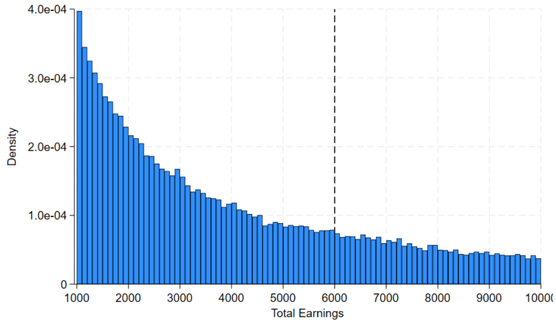

Note: These figures plot the distribution of total earnings prior to separation for workers whose hiring was subsidized by WOTC and who were eligible for WOTC through SNAP receipt. Panel (a) plots this for workers at the 50 firms that most heavily use WOTC. Panel (b) plots this for the WOTC-utilizing firms who are in the top quartile in terms of applying to WOTC for WOTC eligible hires. Panel (c) plots this for workers at employment services firm (i.e. NAICS 5613). Panels (d) and (e) plot this for the years 2007 to 2013 and after 2014 across all firms.

Figure B13: Histograms of total hours worked for WOTC-certified workers

(a) Top 50 WOTC firm

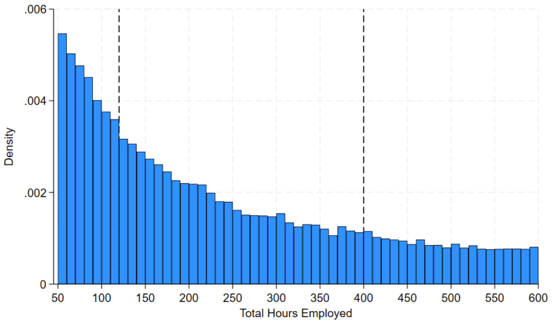

(b) WOTC high success firm

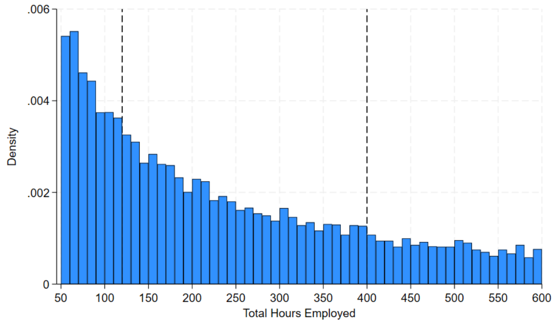

(c) Staffing firm

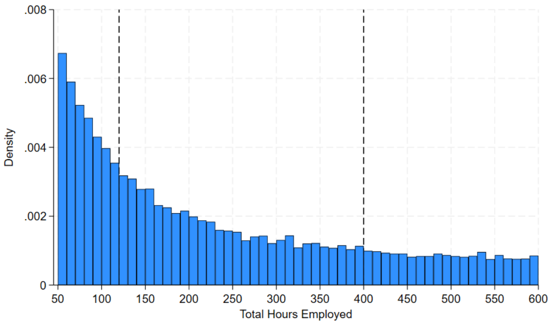

(d) 2007 to 2013

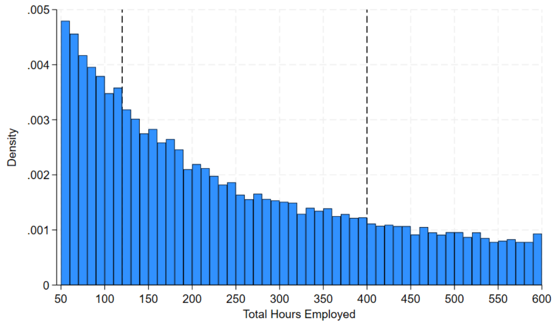

(e) 2014 to 2019

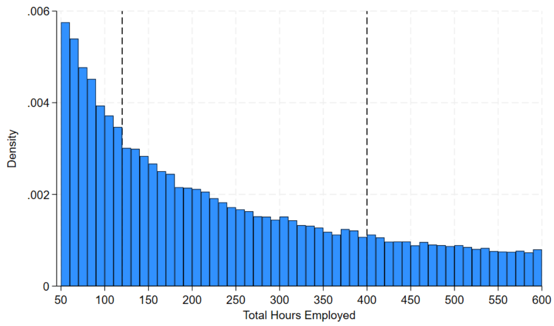

Note: These figures plot the distribution of total hours worked prior to separation for workers whose hiring was subsidized by WOTC and who were eligible for WOTC through SNAP receipt. Panel (a) plots this for workers at the 50 firms that most heavily use WOTC. Panel (b) plots this for the WOTC-utilizing firms who are in the top quartile in terms of applying to WOTC for WOTC eligible hires. Panel (c) plots this for workers at employment services firm (i.e. NAICS 5613). Panels (d) and (e) plot this for the years 2007 to 2013 and after 2014 across all firms.

Figure B14: Aggregate effects

(a) WOTC certifications

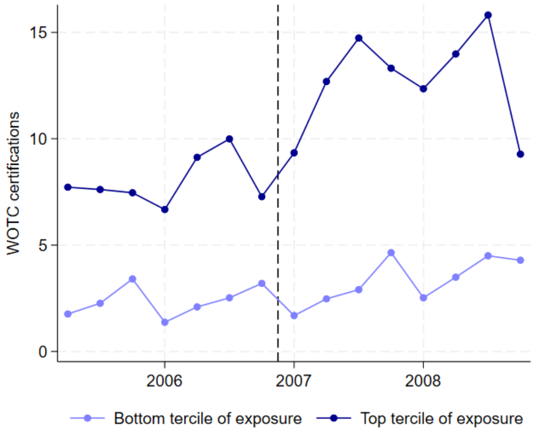

(b) WOTC certifications per hire

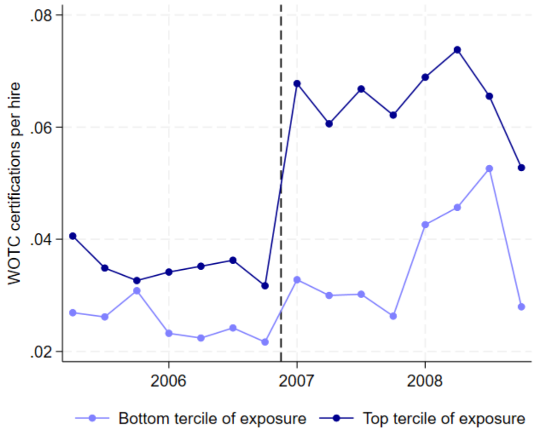

(c) Total hires

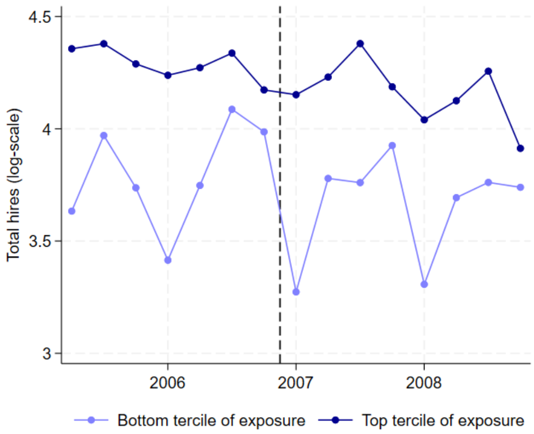

(d) Total payroll ($1,000s)

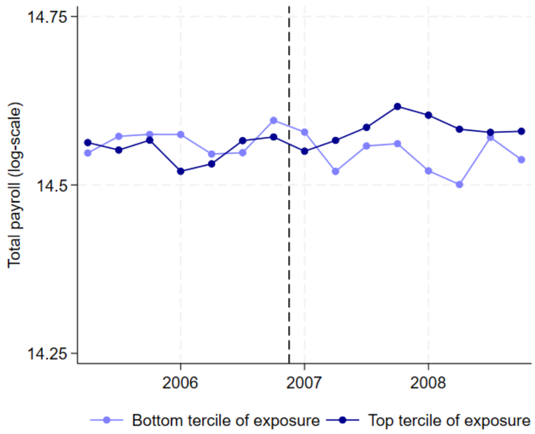

(e) Average starting wage (WOTC)

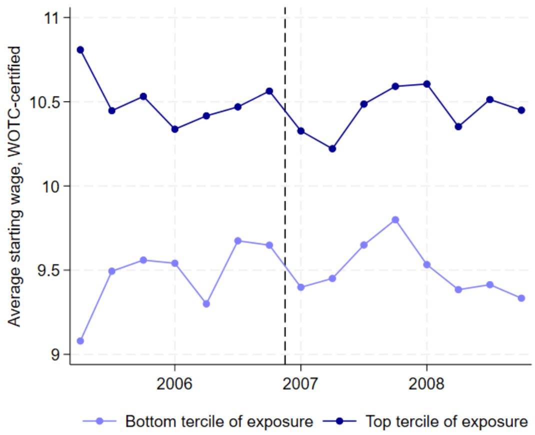

Note: Figures plot outcomes by whether the firm was in the top tercile of exposure relative to the bottom tercile. Bottom tercile firms are reweighted to align their pre-2007 payroll-by-hires distribution with the top tercile firms' distribution. Exposure is defined as the fraction of the industry's hires in the pre-2007 period that received SNAP and were between the ages of 25-39. Sample is firms that had adopted WOTC prior to 2007 and whose exposure fell in the bottom or top tercile of the exposure distribution.

Figure B15: Examples of WOTC information collection on job applications

(a) Example of third-party WOTC screening

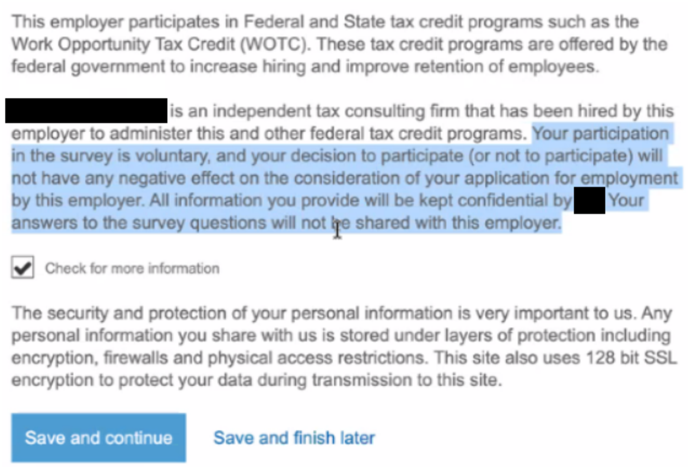

This employer participates in Federal and State tax credit programs such as the Work Opportunity Tax Credit (WOTC). These tax credit programs are offered by the federal government to increase hiring and improve retention of employees.

[redacted] is an independent tax consulting firm that has been hired by this employer to administer this and other federal tax credit programs. Your participation in the survey is voluntary, and your decision to participate (or not to participate) will not have any negative effect on the consideration of your application for employment by this employer. All information you provide will be kept confidential by [redacted] Your answers to the survey questions will not be shared with this employer.

Check for more information

The security and protection of your personal information is very important to us. Any personal information you share with us is stored under layers of protection including encryption, firewalls and physical access restrictions. This site also uses 128 bit SSL encryption to protect your data during transmission to this site.

Save and continue   Save and finish later

(b) Example of direct screening by the company

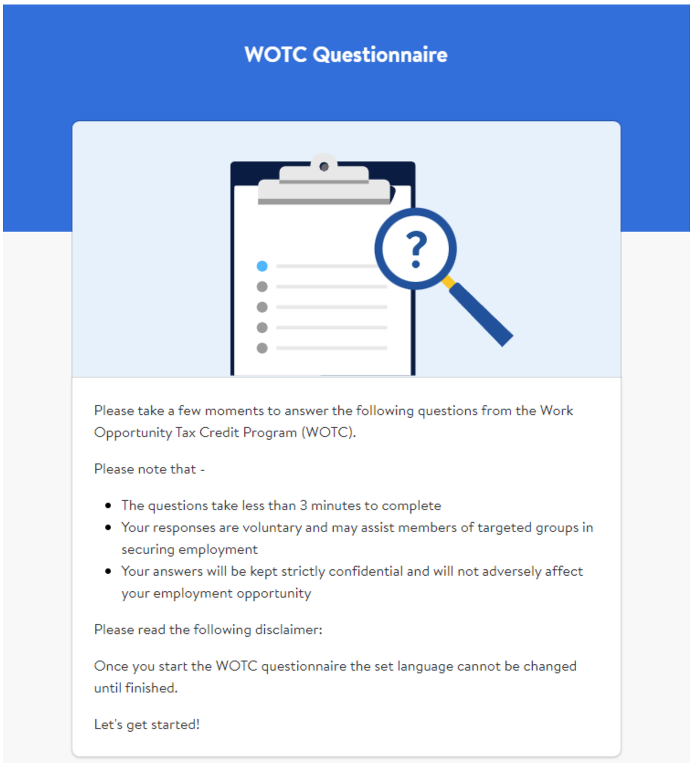

WOTC Questionnaire

Please take a few moments to answer the following questions from the Work Opportunity Tax Credit Program (WOTC).

Please note that -

- The questions take less than 3 minutes to complete
- Your responses are voluntary and may assist members of targeted groups in securing employment
- Your answers will be kept strictly confidential and will not adversely affect your employment opportunity

Please read the following disclaimer:

Once you start the WOTC questionnaire the set language cannot be changed until finished.

Let's get started!

Note: This figure contains images from job applications that include WOTC screening questions.

B.2 Additional Tables

Table B1: Effect of WOTC on individual outcomes, DiD analysis from SNAP expansion (unemployed sample)

	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Panel A: Stacked sample						
Treat x post	−0.0028	−0.0011	−6.9574	0.0115***	14.6411	−0.0001
	(0.0028)	(0.0031)	(8.9215)	(0.0009)	(11.2310)	(0.0010)
Dep var mean	0.1632	0.1932	361.3695	0.0192	988.8316	0.0191
Observations	345,040	345,040	345,040	345,040	345,040	345,040
Panel B: Age 25 cutoff sample						
	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Treat x post	−0.0011	−0.0023	0.2047	0.0152***	5.3376	−0.0007
	(0.0042)	(0.0046)	(12.3641)	(0.0015)	(14.7700)	(0.0014)
Dep var mean	0.2196	0.2564	455.2319	0.0334	1,149.4041	0.0213
Observations	194,184	194,184	194,184	194,184	194,184	194,184
Panel C: Age 40 cutoff sample						
	New hire	Employed	Earnings	Certified (new hire)	Public benefit (amt)	Convicted (=1)
Treat x post	−0.0050	0.0003	−16.1833	0.0068***	26.6256	0.0007
	(0.0036)	(0.0040)	(12.7742)	(0.0007)	(17.2781)	(0.0015)
Dep var mean	0.0945	0.1161	246.9811	0.0019	793.1446	0.0165
Observations	150,856	150,856	150,856	150,856	150,856	150,856

Note: Table reports difference-in-differences estimates of the effect of WOTC from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.2) on new hires, employment, earnings, new hire into a WOTC-certified job, social assistance utilization, and criminal activity. Sample is restricted to individuals who entered quarter unemployed. Controls include monthly age fixed effects and quarter-year by age sample fixed effects in Panel A. Controls include monthly age fixed effects and quarter-year fixed effects in Panels B and C. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B2: Effect of WOTC on employment and hiring by individual characteristics

	Race		Gender		Past conviction	
Panel A: New Hire						
	Non-black	Black	Female	Male	Yes	No
Treat x post	0.0023	0.0008	0.0008	0.0005	0.0013	0.0007
	(0.0037)	(0.0024)	(0.0024)	(0.0038)	(0.0056)	(0.0022)
Dep var mean	0.1591	0.1349	0.1511	0.1190	0.1580	0.1399
Observations	193,648	399,840	441,014	153,148	92,247	501,915
Panel B: Employed						
	Race		Gender		Past conviction	
Black	Female	Male	Yes	No		
Treat x post	0.0134*	0.0041	0.0078	0.0045	0.0192*	0.0050
	(0.0075)	(0.0052)	(0.0050)	(0.0079)	(0.0101)	(0.0047)
Dep var mean	0.5093	0.5211	0.5508	0.4211	0.4484	0.5299
Observations	193,648	399,840	441,014	153,148	92,247	501,915

Note: Table reports difference-in-differences estimates testing the effect of WOTC from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.2) on labor market outcomes by individual characteristics. Panel A examines whether the individual is newly hired, dividing the sample by the individual's race (columns 1 and 2), gender (columns 3 and 4), and having a past criminal conviction (columns 5 and 6). Panel B focuses on the same set of outcomes, but examines whether that the individual is employed rather than newly hired. Controls include monthly age fixed effects and quarter-year by age sample fixed effects. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B3: Effect of WOTC on employment and hiring by firm type

	Panel A: New hire			
	Top 50 firm	High succ firm	Staffing firm	High comp industry
Treat x post	−0.0003	−0.0002	0.0004	−0.0011
	(0.0008)	(0.0007)	(0.0010)	(0.0014)
Dep var mean	0.0236	0.0179	0.0338	0.0653
Observations	594,524	594,524	594,524	594,524
	Panel B: Employed			
	Top 50 firm	High succ firm	Staffing firm	High comp industry
Treat x post	−0.0018	0.0004	0.0006	0.0076**
	(0.0022)	(0.0018)	(0.0020)	(0.0037)
Dep var mean	0.0836	0.0551	0.0773	0.2476
Observations	594,524	594,524	594,524	594,524

Note: Table reports difference-in-differences estimates testing the effect of WOTC from the 2007 expansion to SNAP recipients aged 25-39 in Equation (2.2) on labor market outcomes by firm type. Panel A examines whether the individual is newly hired at one of the 50 firms that are the heaviest users of WOTC (column 1), have a WOTC application success rate above the 75^{th} percentile (column 2), are an employment services firm, i.e. NAICS 5613 (column 3), or are a highly competitive firm (HHI below the 25^{th} percentile). Panel B focuses on the same set of outcomes, but examines whether that the individual is employed rather than newly hired. Controls include monthly age fixed effects and quarter-year by age sample fixed effects. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B4: Balance on pre-determined characteristics, RD analysis

	Demographics		Employed		Earnings	
	(1) Female (=1)	(2) Black (=1)	(3) (t-1)	(4) (t-2)	(5) (t-1)	(6) (t-2)
RD_Estimate	0.0012	-0.0004	-0.0006	0.0002	3.8123	7.8239
	(0.0021)	(0.0021)	(0.0020)	(0.0020)	(11.7246)	(11.3883)
Dep var mean	0.6540	0.3054	0.4833	0.4779	1972.7684	1857.3763
Bandwidth	2.5335	4.1067	3.2818	3.0561	4.3648	4.2236
Effective Obs	2453543	4000381	3045416	2794327	4057170	3878363

Note: Table reports regression discontinuity estimates testing whether pre-determined covariates vary discontinuously across the age threshold. The sample includes both the pre-2007 and post-2007 periods. Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors in parentheses are clustered at the individual level; p-values in square brackets.

Table B5: SNAP usage as a function of WOTC eligibility, RD analysis

	Pre-2007	Post-2007
RD_Estimate	0.0000	0.0009
	(0.0008)	(0.0010)
Dep var mean	1.0000	1.0000
Bandwidth	2.7560	4.1658
Effective Obs	8,154,275	13,003,920

Note: Table reports regression discontinuity estimates testing whether SNAP receipt changes as a function of WOTC eligibility, i.e., whether individuals who are WOTC eligible are more likely to apply for and receive SNAP benefits. The running variable is the distance between the individual's age and the age cutoff for WOTC eligibility. The sample includes all individuals in the Wisconsin Administrative Data Core from 1998 to 2019. Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors in parentheses are clustered at the individual level; p-values in square brackets.

Table B6: Effect of WOTC on individual outcomes for the age 25 discontinuity (pre-2007)

	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
Panel A: Full sample						
RD_Estimate	−0.0008	0.0012	4.4106	−0.0060***	2.3184	−0.0006
	(0.0024)	(0.0039)	(21.3345)	(0.0019)	(7.6854)	(0.0010)
Dep var mean	0.193	0.570	2,164.937	0.007	1,006.151	0.015
Bandwidth	2.609	2.198	2.469	1.193	2.682	2.308
Effective Obs	670,323	564,441	628,505	95,351	690,768	331,729
Panel B: Unemployed sample						
	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
RD_Estimate	−0.0002	−0.0009	−2.0195	−0.0086***	3.3656	−0.0010
	(0.0031)	(0.0033)	(9.5933)	(0.0024)	(10.7880)	(0.0016)
Dep var mean	0.222	0.254	508.083	0.008	1,197.885	0.020
Bandwidth	3.138	2.909	2.947	1.548	3.101	2.281
Effective Obs	461,858	426,329	438,208	65,655	461,858	182,393

Note: Table reports regression discontinuity estimates for employment, earnings, program utilization, and criminal activity in the pre-2007 periods. The running variable is distance between the individual's age and the age cutoff for WOTC eligibility (age 25 before 2007 and age 40 in 2007 and afterwards). Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B7: Effect of WOTC on individual outcomes for the age 40 discontinuity (post-2007)

	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
Panel A: Full sample						
RD_Estimate	−0.0007	−0.0001	−6.7245	−0.0102***	−2.4194	0.0000
	(0.0010)	(0.0022)	(14.8457)	(0.0004)	(5.0939)	(0.0003)
Dep var mean	0.097	0.422	1,986.687	0.003	1,031.352	0.010
Bandwidth	3.119	4.205	4.051	1.904	3.629	5.910
Effective Obs	2,237,715	3,051,764	2,926,116	1,310,450	2,612,025	4,333,731
Panel B: Unemployed sample						
	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
RD_Estimate	0.0003	0.0012	5.1002	−0.0108***	−4.2548	−0.0001
	(0.0013)	(0.0015)	(4.9809)	(0.0004)	(6.8718)	(0.0004)
Dep var mean	0.092	0.116	269.716	0.003	1,111.432	0.013
Bandwidth	2.962	2.809	3.038	2.585	3.704	5.094
Effective Obs	1,316,675	1,239,692	1,355,817	1,162,364	1,666,465	2,338,938

Note: Table reports regression discontinuity estimates for employment, earnings, program utilization, and criminal activity in the post-2007 periods. The running variable is distance between the individual's age and the age cutoff for WOTC eligibility (age 25 before 2007 and age 40 in 2007 and afterwards). Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B8: Effect of WOTC on employment and hiring by individual characteristics

	Race		Gender		Past conviction	
Panel A: New Hire						
	Non-black	Black	Female	Male	Yes	No
RD_Estimate	0.0001	−0.0011	0.0013	−0.0020	0.0012	0.0001
	(0.0035)	(0.0023)	(0.0025)	(0.0030)	(0.0042)	(0.0022)
Dep var mean	0.465	0.465	0.500	0.394	0.356	0.492
Bandwidth	3.653	3.728	3.035	5.359	4.284	2.843
Effective Obs	1,077,393	2,490,163	1,953,968	1,793,276	756,916	2,284,982
Panel B: New hire						
	Race		Gender		Past conviction	
	Non-black	Black	Female	Male	Yes	No
RD_Estimate	−0.0019	−0.0001	−0.0001	−0.0017	−0.0022	−0.0005
	(0.0019)	(0.0012)	(0.0013)	(0.0017)	(0.0020)	(0.0011)
Dep var mean	0.145	0.114	0.131	0.110	0.115	0.125
Bandwidth	3.068	2.726	2.777	2.858	4.113	2.826
Effective Obs	902,147	1,810,147	1,791,604	937,173	725,794	2,217,935

Note: Table reports regression discontinuity estimates for labor market outcomes by individual characteristics. Panel A examines whether the individual is newly hired, dividing the sample by the individual's race (columns 1 and 2), gender (columns 3 and 4), and having a past criminal conviction (columns 5 and 6). Panel B focusses on the same set of outcomes, but examines whether that the individual is employed rather than newly hired. The running variable is distance between the individual's age and the age cutoff for WOTC eligibility (age 25 before 2007 and age 40 in 2007 and afterwards). Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level. Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B9: Effect of WOTC on employment and hiring by firm type and year

	Top 50 firm	High succ firm	Staffing firm	High comp industry	1998-2006	2007-2013	2014-2019
Panel A: Employed							
	Firm-level				Years		
RD_Estimate	0.0000	0.0004	0.0003	−0.0013	−0.0010	0.0022	−0.0017
	(0.0009)	(0.0007)	(0.0009)	(0.0016)	(0.0040)	(0.0031)	(0.0032)
Dep var mean	0.064	0.035	0.070	0.217	0.577	0.401	0.442
Bandwidth	3.686	3.576	3.832	4.002	2.310	4.399	4.550
Effective Obs	3,601,745	3,438,705	3,684,042	3,929,298	535,047	1,613,972	1,628,347
Panel B: New hire							
	Firm-level				Years		
	Top 50 firm	High succ firm	Staffing firm	High comp industry	1998-2006	2007-2013	2014-2019
RD_Estimate	0.0002	0.0001	0.0002	0.0001	−0.0013	0.0013	−0.0028*
	(0.0004)	(0.0003)	(0.0005)	(0.0006)	(0.0024)	(0.0012)	(0.0015)
Dep var mean	0.018	0.010	0.028	0.051	0.182	0.086	0.107
Bandwidth	2.296	2.609	3.168	3.646	2.695	4.145	3.637
Effective Obs	2,204,601	2,535,660	3,110,432	3,520,309	630,557	1,519,191	1,281,592

Note: Table reports regression discontinuity estimates for employment (panel A) and being hired (panel B) at firms that might plausibly be more responsive to WOTC and over different time intervals. The running variable is distance between the individual's age and the age cutoff for WOTC eligibility (age 25 before 2007 and age 40 in 2007 and afterwards). The first four columns examines the outcomes of employment at the 50 firms that have the most WOTC subsidized jobs, firms with a high success rate in applying for WOTC for their WOTC-eligible hires, staffing services agencies (NACIS 5613), and firms in more competitive industries (as measured by having a Herfindahl index in the bottom 25% of industries, using data from the Economic Census). Columns 5 to 7 measure the outcomes within distinct time intervals (1997 to 2006, 2007 to 2013, and 2014 to 2019). Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B10: Placebo RD analysis

	New hire	Employed	Earnings	Certified	Public benefit (amt)	Convicted (=1)
RD_Estimate	−0.0002	−0.0021	−11.5450	−0.0007*	−2.0942	0.0002
	(0.0010)	(0.0017)	(10.0275)	(0.0004)	(3.6145)	(0.0003)
Dep var mean	0.1560	0.5452	2,194.8615	0.0258	1,027.5926	0.0148
Bandwidth	2.5791	2.5680	2.6078	3.1502	2.5395	3.7762
Effective Obs	3,030,117	3,030,117	3,126,338	3,272,132	3,030,117	4,101,431

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Note: Table reports the relationship between individual-level outcomes and the difference between the individual's age and a placebo WOTC age eligibility threshold. For the pre-2007 period, the placebo age eligibility threshold is age 40, while for the post-2007 period, the placebo age eligibility threshold is age 25. This checks whether there is a relationship between the age threshold and the individual level outcomes during the years when this age threshold is not relevant for WOTC eligibility. Panel A examines outcomes for the pre-2007 period, while Panel B focuses on the post-2007 period. Each specification uses a linear polynomial, triangular kernel, and MSE-optimal bandwidth estimated following Calonico et al. (2017). Standard errors are clustered at the individual level.

Table B11: Effect of WOTC expiration on the hazard rate of exit from employment

	Hazard model	Hazard model	Hazard model
Time	-0.0028***	-0.0014***	-0.0010***
	(0.0000)	(0.0000)	(0.0000)
Post	14.9490***	5.6687***	3.1156***
	(0.1283)	(0.0447)	(0.0245)
WOTC	-0.1431	-0.0230	0.0041
	(0.1999)	(0.0608)	(0.0276)
Post X WOTC	0.1215	-0.0408	0.0024
	(0.2563)	(0.0893)	(0.0491)
Post X Time	-0.0023***	-0.0007***	-0.0003***
	(0.0000)	(0.0000)	(0.0000)
Post X WOTC X Time	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)
WOTC X Time	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)
Lower bound	5,000	4,000	3,000
Upper bound	7,000	8,000	9,000
Observations	696,590	779,176	885,725

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Note: Table estimates a hazard model of employment, testing whether the hazard rate of separation increase for WOTC workers relative to non-WOTC workers around \$6000 in earnings, when the WOTC subsidy expires. We restrict to employment stints in 2007 and after (when the WOTC application micro-data is higher quality) at firms who submitted at least one WOTC application in the quarter of hire. We also restrict to workers aged 18 to 39 at time of hiring to maintain comparability of age with WOTC hires. The full model regresses the hazard rate of separation on a dummy variable for if the stint is associated with a WOTC application (WOTC), dollar of earnings (*earnings*), a dummy variable for if the earnings level is above \$6000, and the full set of interactions of these variables ($postXearnings, WOTCXearnings, postXWOTC$, and $postXWOTCXearnings$). We are interested in whether the hazard rate changes discontinuously around \$6000, i.e., the coefficient on $postXWOTC$. For robustness, we estimate this within three different bandwidths around \$6000.

Table B12: Substitutability of WOTC-eligible and ineligible SNAP beneficiaries

	Eligible hires	Eligible hires	Ineligible hires	Ineligible hires
Eligible exits	0.124***	0.084***	0.107***	0.086***
	(0.025)	(0.013)	(0.020)	(0.011)
Ineligible exits	0.151***	0.114***	0.119***	0.087***
	(0.026)	(0.014)	(0.027)	(0.013)
Firm FEs	Yes	Yes	Yes	Yes
Year-Quarter FEs	Yes	Yes	Yes	Yes
Total Exits FEs	No	Yes	No	Yes
Observations	615,668	615,264	615,668	615,264

Note: Table reports estimates of how many hires a firm made in a quarter as a function of separations from the firm in that quarter. Column 1 examines the number of hires of individuals who are between ages 37-39 and received SNAP benefits for 3 of the past 5 months (eligible for WOTC), while column 2 examines hires of those between ages of 42-44 and received SNAP benefits for 3 of the past 5 months (ineligible for WOTC). In both cases, the dependent variable is regressed on the number of exits from the firm in that quarter of individuals with those characteristics. All specifications include fixed effects for firm and year-quarter, as well as the total number of exits from the firm in that quarter. Standard errors are clustered at the firm level.

Table B13: Aggregate firm-level effects of WOTC

	WOTC certs	WOTC certs/hire	Hires (logs)	Payroll (logs)	Starting wage (WOTC)
Top tercile	5.60***	0.0100*	0.49***	−0.013	1.04***
	(1.73)	(0.0056)	(0.16)	(0.11)	(0.22)
Top tercile x post	3.78**	0.019***	0.016	0.052	−0.093
	(1.53)	(0.0059)	(0.068)	(0.045)	(0.16)
Mean, bottom tercile	3.275	0.0300	3.639	14.92	9.874
Observations	4,592	4,512	4,512	4,592	2,752

Note: Table reports estimates from a firm-level regression of the outcome denoted by the column header on whether the firm was in the top tercile of exposure relative to the bottom tercile, year-quarter fixed effects, and the interaction between top tercile and whether the year-quarter was after the policy change. Bottom tercile firms are rewighted to align their pre-2007 payroll-by-hires distribution with the top tercile firms' distribution. Exposure is defined as the fraction of the industry's hires in the pre-2007 period who received SNAP and were between the ages of 25-39. Sample is firms that had adopted WOTC prior to 2007 and whose exposure fell in the bottom or top tercile of the exposure distribution. Standard errors clustered at the firm level.

Table B14: Audit study analysis

	Type of firm		Diff-in-RD	Applicant type (WOTC firms)	
	WOTC firm	Non-WOTC firm	.	Black	Non-college
RD_Estimate	−0.038	0.011	−0.049	−0.054	−0.026
	(0.029)	(0.038)	(0.048)	(0.034)	(0.032)
Dep var mean	0.241	0.242		0.241	0.242
Bandwidth	7.627	8.070		6.657	7.547
Effective Obs	19,946	12,755		8,627	10,227

Note: Table reports estimates from audit study.

Appendix C: Chapter 3

C.1 Data Appendix

This section details the process by which I fill in the two-digit occupation codes of sons and fathers. I discuss how I construct occupation codes for resident sons, followed by non-resident sons, and finally, I describe how I construct the occupation code for fathers.

C.1.1 Occupation Code of Resident Sons

The IHDS contains information on the sources of household income from joint production activities (income from cultivation, animal husbandry, and business),[1], each resident household member's participation in each of these activities, and their level of participation (the number of hours worked in a day and the number of days worked in a year).[2] It also asks every resident household member if they engaged in any agricultural wage labour, non-agricultural wage labour, or salaried work and if yes, for how many hours they worked in a year. As many individuals in my data set worked in more than one work activity in the year of the survey, I assign their main work activity (cultivation, working in business 1, working in business 2, working in business 3, salaried/agricultural wage labourer/non-agricultural wage labourer, or animal husbandry) as the activity in which they work the highest hours per year, provided that number is greater than 240 hours/year.[3]

The IHDS also provides two-digit occupation codes for salaried work/agricultural wage labour/non-agricultural wage labour (the variable called "WS4" contained in the individual

[1] The IHDS reports information on a maximum of 3 businesses owned by a household.

[2] Information on the number of hours/days worked in animal husbandry is not available; instead, the IHDS asks if the member takes care of the animals "usually", "sometimes", or "never".

[3] I assign animal husbandry as the main work activity if the maximum hours worked per year in all other work activities is less than 240 and the household member "usually" takes care of animals.

file) and each of the possibly three businesses of the household (variables "NF1B", "NF21B", and "NF41B", respectively, contained in the household file). Thus, if the main work activity is classified as working in business 1 (2) (3), then I fill in the occupation code for this resident by the occupation code reported in the variable "NF1B" ("NF21B") ("NF41B"). If the main work activity is classified as work/agricultural wage labour/non-agricultural wage labour, I disaggregate workers using the variable "WS4". Lastly, I impute the occupation code as 61 if the main work activity of that resident is cultivation or animal husbandry.

I fill in some of the missing information of the occupation codes[4] from the variable "RO7" that reports the "primary occupation status"[5] of each resident. If the residents report their primary occupation as cultivation or allied agriculture, I classify their (originally missing) occupation code as 61 (cultivators) and if they report their primary occupation as agricultural labour, I classify their (originally missing) occupation code as 63.

I do not use this variable as the first step in deducing the main work activity as there are many cases in which the number of hours reported in a different work activity is higher than the one stated as the "primary occupation". Azam (2013) and Kundu and Sen (2021) also do not rely on the "primary occupation status" of the residents, but instead use the indicators provided in the IHDS for different work activities (cultivation, animal care, business, salaried, agricultural wage labour, non-agricultural wage labour), each of which is reported as "none" (i.e. 0 hours worked in that activity), missing hours, less than 240 hours/year, part time (greater than 240 hours/year, but not full time), or full time (greater than 250 days/year and more than 2000 hours/year).[6] There are a few cases in which a resident reports having two full time work activities, and there are many cases in which an individual reports having no full time activity but multiple part-time activities. Thus, Azam (2013) and Kundu and Sen (2021) end up using a main-work-activity-classification procedure that relies on

[4]This information could be missing, for example, if the main work activity is missing in cases where the individual did not report working hours and/ or days, is unemployed, retired, does household work, or is a student.

[5]The primary occupation status of each resident takes one of the following values: Cultivation, Allied Agriculture, Agricultural Wage Labour, Non-Agricultural Wage labour, Artisan/Independent Work, Small Business, Organized Business, Professional, Salaried, Retired, Housework, Student, Unemployed, Too young/unfit, and Others

[6]Since the number of hours worked in animal care is not reported, the indicator for animal work takes the following values: none, less than 240 hours (i.e. work "sometimes" in animal care), and part time (i.e. work "usually" in animal care)

a subjective order in which they consider these indicators (for example, they first try to categorize residents as salaried/wage labourers if they spend more than 240 hours/year in salaried/wage work, ignoring the possibility that they could indeed be classified as full time in some other work activity. They then try to classify residents as working in a business, and so on. The advantage of using the highest number of hours worked to deduce the main work activity is that I do not have to use an arbitrary order to go through these indicators and I can also infer the particular business in which the resident spends the maximum of his working time, so as to fill in the occupation code of that specific business, if needed.

C.1.2 Occupation Code of Non-Resident Sons

Filling in the two-digit occupation codes of non-resident sons is a more straight-forward process as the variable "NR11" in the non-resident file of IHDS-II directly contains their occupation code.

C.1.3 Occupation Code of Fathers

I fill in the two-digit occupation codes of fathers in the following way. More than 50% of sons in my sample are also household heads or husbands of household heads. The household file of IHDS-II contains a variable called "ID18A", which reports the two-digit occupation code of the father of a male household head or the husband's father of a female household head, even if they are dead or are not considered a part of the same household. Thus, I can directly obtain information on father's occupation choice for male household heads and husbands of female household heads from this variable. For sons who are not household heads and have resident fathers, I fill in the father's occupation code in a manner similar to the one I use for resident sons i.e. I infer their main work activity using information on hours worked by the father in different work activities, fill in the occupation codes using the "WS4", "NF1B", "NF21B", or "NF41B" variables corresponding to the father's row in the individual file of the IHDS, wherever appropriate, and replace some missing information in the father's occupation code using information reported in the "RO7" variable corresponding to the father's row. For sons who are not household heads and have non-resident fathers, I use information contained in the "NR11" variable in the non-resident file that corresponds to the individual who is the father of the son under consideration.

Table C1: NCO-1968 Two-Digit Occupation Codes

Code	Occupation
	PROFESSIONAL, TECHNICAL, AND RELATED WORKERS
00	Physical Scientists
01	Physical Science Technicians
02	Architects, Engineers, Technologists, and Surveyors
03	Engineering Technicians
04	Aircraft and Ships Officers
05	Life Scientists
06	Life Science Technicians
07	Physicians and Surgeons (Allopathic Dental and Veterinary Surgeons)
08	Nursing and Other Medical and Health Technicians
09	Scientific, Medical, and Technical Persons, Other
10	Mathematicians, Statisticians, and Related Workers
11	Economists and Related Workers
12	Accountants, Auditors, and Related Workers
13	Social Scientists and Related Workers
14	Jurists
15	Teachers
16	Poets, Authors, Journalists, and Related Workers
17	Sculptors, Painters, Photographers, and Related Creative Artists
18	Composers and Performing Artists
19	Professional Workers, n.e.c.
	ADMINISTRATIVE, EXECUTIVE, AND MANAGERIAL WORKERS
20	Elected and Legislative Officials
21	Administrative and Executive Officials, Government and Local Bodies
22	Working Proprietors, Directors and Managers, Wholesale and Retail Trade
23	Directors and Managers, Financial Institutions
24	Working Proprietors, Directors and Managers Mining, Construction, Manufacturing, and Related Concerns
25	Working Proprietors, Directors, Managers and Related Executives, Transport, Storage, and Communication
26	Working Proprietors, Directors and Managers, Other Service
29	Administrative, Executive and Managerial Workers, n.e.c.
	CLERICAL AND RELATED WORKERS
30	Clerical and Other Supervisors
31	Village Officials
32	Stenographers, Typists, and Card and Tape Punching Operators
33	Book-keepers, Cashiers, and Related Workers
34	Computing Machine Operators
35	Clerical and Related Workers, n.e.c.
36	Transport and Communication Supervisors
37	Transport Conductors and Guards
38	Mail Distributors and Related Workers
39	Telephone and Telegraph Operators

Continued on next page

Table C1 – *Continued from previous page*

Code	Occupation
	SALES WORKERS
40	Merchants and Shopkeepers, Wholesale and Retail Trade
41	Manufacturers, Agents
42	Technical Salesmen and Commercial Travellers
43	Salesmen, Shop Assistants, and Related Workers
44	Insurance, Real Estate, Securities, and Business Service Salesmen and Auctioneers
45	Money Lenders and Pawn Brokers
49	Sales Workers, n.e.c.
	SERVICE WORKERS
50	Hotel and Restaurant Keepers
51	House Keepers, Matron, and Stewards (Domestic and Institutional)
52	Cooks, Waiters, Bartenders, and Related Worker (Domestic and Institutional)
53	Maids and Other House Keeping Service Workers n.e.c.
54	Building Caretakers, Sweepers, Cleaners, and Related Workers
55	Launderers, Dry-cleaners, and Pressers
56	Hair Dressers, Barbers, Beauticians, and Related Workers
57	Protective Service Workers
59	Service Workers, n.e.c.
	FARMERS, FISHERMEN, HUNTERS, LOGGERS, AND RELATED WORKERS
60	Farm Plantation, Dairy and Other Managers and Supervisors
61	Cultivators
62	Farmers other than Cultivators
63	Agricultural Labourers
64	Plantation Labourers and Related Workers
65	Other Farm Workers
66	Forestry Workers
67	Hunters and Related Workers
68	Fishermen and Related Workers
PRODUCTION AND RELATED WORKERS, TRANSPORT EQUIPMENT OPERATORS AND LABOURERS	
71	Miners, Quarrymen, Well Drillers, and Related Workers
72	Metal Processors
73	Wood Preparation Workers and Paper Makers
74	Chemical Processors and Related Workers
75	Spinners, Weavers, Knitters, Dyers, and Related Workers
76	Tanners, Fellmongers, and Pelt Dressers
77	Food and Beverage Processors
78	Tobacco Preparers and Tobacco Product Makers
79	Tailors, Dress Makers, Sewers, Upholsterers, and Related Workers
80	Shoe Makers and Leather Goods Makers
81	Carpenters, Cabinet and Related Wood Workers
82	Stone Cutters and Carvers
83	Blacksmiths, Tool Makers, and Machine Tool Operators

Continued on next page

Table C1 – *Continued from previous page*

Code	Occupation
84	Machinery Fitters, Machine Assemblers, and Precision Instrument Makers (except Electrical)
85	Electrical Fitters and Related Electrical and Electronic Workers
86	Broadcasting Station and Sound Equipment Operators and Cinema Projectionists
87	Plumbers, Welders, Sheet Metal, and Structural Metal Preparers and Erectors
88	Jewellery and Precious Metal Workers and Metal Engravers (Except Printing)
89	Glass Formers, Potters, and Related Workers
90	Rubber and Plastic Product Makers
91	Paper and Paper Board Products Makers
92	Printing and Related Workers
93	Painters
94	Production and Related Workers, n.e.c.
95	Bricklayers and Other Constructions Workers
96	Stationery Engines and Related Equipment Operators, Oilers and Greasers
97	Material Handling and Related Equipment Operators, Loaders and Unloaders
98	Transport Equipment Operators
99	Labourers, n.e.c.
	WORKERS NOT CLASSIFIED BY OCCUPATIONS
X0	New Workers Seeking Employment
X1	Workers Reporting Occupations Unidentifiable
X9	Workers Not Reporting Any Occupation
AA	Housewife/Household work
BB	Student/Too Young to Work
CC	Retired/Too Old to Work
DD	Disabled/Unfit to Work
EE	Out of Labour Force n.e.c

Table C2: Occupation Categories in Azam (2013)

	Category	NCO-1968 Occupation Code
1	White collar	0-29
2	Skilled	30-45, 49, 71-98
3	Unskilled	50-59, 63-68, 99
4	Farmers	60-62